arXiv:1702.04551v1 [cs.LO] 15 Feb 2017

A Logical Study of Some Common Principles of Inductive Definition and its Implications for Knowledge Representation

Marc Denecker, Bart Bogaerts, Joost Vennekens

February 16, 2017

Abstract

The definition is a common form of human expert knowledge, a building block of formal science and mathematics, a foundation for database theory and is supported in various forms in many knowledge representation and formal specification languages and systems. This paper is a formal study of some of the most common forms of inductive definitions found in scientific text: monotone inductive definition, definition by induction over a well-founded order and iterated inductive definitions. We define a logic of definitions offering a uniform formal syntax to express definitions of the different sorts, and we define its semantics by a faithful formalization of the induction process. Several fundamental properties of definition by induction emerge: the non-determinism of the induction process, the confluence of induction processes, the role of the induction order and its relation to the inductive rules, how the induction order constrains the induction process and, ultimately, that the induction order is irrelevant: the defined set does not depend on the induction order. We propose an inductive construction capable of constructing the defined set *without* using the induction order. We investigate borderline definitions of the sort that appears in definitional paradoxes.

1 Introduction

This paper is a formal scientific study of certain types of definitions as they appear in mathematical and scientific text. The *definition* is one of the building blocks of science and mathematics and its use is ubiquitous there. Consequently, it received due attention from logicians and computer scientists. Inductive definitions were investigated in metamathematical studies (Moschovakis, 1974a; Aczel, 1977; Feferman, 1970; Martin-Löf, 1971; Buchholz et al., 1981). Definitions play an important role in many declarative paradigms and systems. In databases (SQL, Datalog), a query is essentially a symbolic definition of a set that the user wants to be calculated. As such, definitions and definability are key concepts in database theory (Abiteboul et al., 1995). Fixpoint logics (Gurevich and Shelah, 1985) have their origin in metamathematical studies of inductive definitions and inductive definability. In logic programming, definitions play an important role as one of the solutions to the problem of explaining the meaning of

logic programs with negation as failure (Clark, 1978; Schlipf, 1995b; Denecker, 1998). In knowledge representation, it is widely recognized that definitions are an important form of human expert knowledge that should be supported in knowledge representation and specification logics (Brachman et al., 1983; Denecker and Ternovska, 2008, 2007). Many declarative systems in various fields of computational logic support some form of definitions, e.g., Minizinc (Nethercote et al., 2007), ProB (Leuschel and Butler, 2008), IDP (De Cat et al., 2016).

The study in this paper focuses on definitions that are, or can be, formulated as a set of informal base rules and inductive rules, possibly equipped with an induction order. This covers the class of non-recursive definitions as a trivial case, but the focus is on induction, evidently. A prototypical example of a definition of this kind is the well-known definition of the transitive closure of a graph.

Definition 1.1. The *reachability graph* $\mathcal{R}$ of a directed graph $\mathcal{G}$ is defined inductively:

- $(d, e) \in \mathcal{R}$ if $(d, e) \in \mathcal{G}$;
- $(d, e) \in \mathcal{R}$ if there exists a vertex f such that $(d, f), (f, e) \in \mathcal{R}$.

An equally well-known definition is the one of the satisfaction relation of propositional logic:

Definition 1.2. Given a propositional vocabulary Σ, the *satisfaction relation* $\models$ between Σ-structures and Σ-formulas of propositional logic is defined by induction over the structure of formulas:

- $I \models P$ if P is a propositional symbol and $P \in I$.
- $I \models \alpha \wedge \beta$ if $I \models \alpha$ and $I \models \beta$.
- $I \models \alpha \vee \beta$ if $I \models \alpha$ or $I \models \beta$ (or both).
- $I \models \neg\alpha$ if $I \not\models \alpha$.

These two definitions are instances of what are probably the two most common forms of inductive definitions in mathematical and formal scientific text. Definition 1.1 is an example of a *monotone inductive definition*. Such definitions were studied extensively in mathematical logic (Moschovakis, 1974a; Aczel, 1977). Definition 1.2 is a definition *by structural induction*, or *by induction on the complexity of the formula*. It is an example of a *definition by induction over a well-founded induction order*: here the induction order is the subformula order. Definitions over an induction order may be non-monotone. For instance, in the fourth rule of Definition 1.2, the condition "$I \not\models \alpha$" is a non-monotone condition, in words "*it is not the case that* $I \models \alpha$". This sort of definition has not been studied so well.[1] These two informal definitions are clear instances of the sort of definitions that we want to study here in this paper and they will serve as running examples throughout the paper. Our study includes also *iterated inductive definitions*, definitions that combine features of monotone induction and induction over

[1] Some reserve the term *inductive definition* for what we call here "monotone inductive definitions", and *recursive definition* for what we call "definitions by induction over an induction order".

a well-founded order. Such definitions were studied in (Feferman, 1970; Martin-Löf, 1971; Buchholz et al., 1981). They are discussed later in the paper. These are the forms of definitions that we study here, and we study them from a logical and semantical point of view. It may seem unlikely that about such common and fundamental objects of mathematical reasoning much remains to be discovered at the semantical level, and yet we shall argue that this is the case.

Definitions in mathematical or scientific text serve to define formal objects, but they are not formal objects themselves. As such, we will refer to them as *informal definitions*. Definitions are propositions that state a particular sort of logical relationship: they define one set (or possibly more than one) *in terms of* other sets which we call the *parameters* of the definition. For instance, the defined set of Definition 1.1 is the reachability graph $\mathcal{R}$ and its unique parameter the graph $\mathcal{G}$; the defined set of Definition 1.2 is the satisfaction relation between Σ-structures and formulas over Σ and the parameter is the vocabulary Σ.

Despite their informal nature, inductive definition found in mathematical text strike us for their precision. The set defined by such an informal definition can often be characterized in two quite different ways: "non-constructively", as the least set closed under rule application, and "constructively", as the set obtained by iterated rule application. By Tarski's least fixpoint theorem, both sets coincide.

Tarski's result, however, holds only for monotone operators. The operator induced by Definition 1.2 is non-monotone due to its fourth rule, and Tarski's theorem does not apply to it: the satisfaction relation $\models$ is **not** the least relation satisfying the rules of Definition 1.2. The least relation does not exist; there are infinitely many minimal sets that are closed under these rules and some are just weird (we return to this in Example 3.12). While experts are aware of this, this comes as a surprise to many people, even those skilled in mathematics. This shows that theoretical understanding of this sort of definition is less widely spread than deserved. What it also shows is that the constructive principle is the more fundamental of the two principles. We choose this principle as the foundation of our study. As such, the sort of definition studied here defines a set by describing how to construct it through an *induction process*. The induction process starts from the empty set and proceeds by applying rules until the set is closed (saturated) under rule application. In case of an induction order, rules must be applied "along" the specified order. The considered class of definitions covers non-inductive definitions as a trivial case (no inductive rules) and also the above sorts of inductive definitions.

Our study is a formal, logical, semantical study of the selected sort of informal definitions. Syntactically, a formal definition will be defined as a set of formal rules

$$\forall \bar{x}\ (P(\bar{t}) \leftarrow \phi)$$

where $P(\bar{t})$ is an atomic formula (the definiendum) with P the defined set or relation and ϕ is a formula (the definiens) of first-order logic (FO). Given a suitable first-order vocabulary Σ to express the concepts of the informal definition, we will say that a formal definition *faithfully expresses* an informal rule-based definition if there is a one-to-one correspondence between formal and informal rules such that the definiendum (i.e., the head) of the rule correctly formalizes the conclusion of the informal rule and

the definiens (i.e., the body) of the formal rule correctly formalizes that of the informal rule. For the running example Definition 1.1, the formalization may be as follows:

$$\Delta_{TC} = \left\{ \begin{array}{l} \forall x \forall y (R(x,y) \leftarrow G(x,y)) \\ \forall x \forall y (R(x,y) \leftarrow \exists z (R(x,z) \wedge R(z,y)) \end{array} \right\}$$

To formalize Definition 1.2, we use the symbol $Sat(i, f)$ to express that f is a formula satisfied in structure i, $Atom(p)$ that p is a propositional atom of the vocabulary and $In(p, i)$ that p is true in structure i, and function symbols $And/2, Or/2, Not/1$ on formulas to express connectives (see Example 3.6 in Section 3 for details).

$$\Delta_{\models} = \left\{ \begin{array}{ll} \forall i \forall p (Sat(i,p) & \leftarrow Atom(p) \wedge In(p,i)) \\ \forall i \forall f \forall g (Sat(i, And(f,g)) & \leftarrow Sat(i,f) \wedge Sat(i,g)) \\ \forall i \forall f \forall g (Sat(i, Or(f,g)) & \leftarrow Sat(i,f) \vee Sat(i,g)) \\ \forall i \forall f (Sat(i, Not(f)) & \leftarrow \neg Sat(i,f)) \end{array} \right\}$$

It can be seen (and it will be shown in a precise mathematical way) that both formal definitions faithfully express the corresponding informal deinition. E.g., in $\Delta_{\models}$, the conclusion $I \models \neg\alpha$ of the fourth informal rule is correctly translated into $Sat(i, Not(f))$ and its condition $I \not\models \alpha$ is faithfully translated into the formula $\neg Sat(i, f)$ (where i stands for I and f for α).

On the semantical level, we will define a model semantics for this formalism. The key concept in this semantics is the formalization of the induction process. For a formal definition Δ, it will be formalized as an increasing, possibly transfinite sequence of sets (or, more generally, of structures)

$$\langle \mathfrak{A}_0, \mathfrak{A}_1, \mathfrak{A}_2, \dots \rangle$$

where $\mathfrak{A}_0 = \emptyset$ and at each stage i a set of applicable rule instances of Δ is applied to obtain $\mathfrak{A}_{i+1}$. The defined set is then the limit of the induction process. A model of a definition will be defined as a structure in which the interpretation of the defined symbol is this defined set.

E.g., a small segment of an induction process for Definition 1.2 in the context of formulas and structures of the vocabulary $\Sigma = \{P, Q\}$ is

$$\emptyset \rightarrow \{(\{P,Q\}, P), (\{P,Q\}, Q)\} \rightarrow \{(\{P,Q\}, P \wedge Q), (\{P,Q\}, P), (\{P,Q\}, Q)\} \rightarrow \dots$$

It is obtained by applying, first, two instances of the base rule of $\Delta_{\models}$ to derive satisfaction of P and Q in the structure $\{P, Q\}$; second, an instance of the rule for conjunctive formulas to derive $P \wedge Q$ in this same structure.

The above formal notion of the induction process is a faithful formalization of the iterated rule application that is inherent to the class of informal definitions that we study here. The concept is a generalization of the formal notion of induction process found in the standard studies of monotone induction such as by Moschovakis (1974a) and Aczel (1977). There, the induction process is formalized as the sequence

$$\langle \emptyset, \Gamma(\emptyset), \Gamma^2(\emptyset), \dots, \Gamma^n(\emptyset), \dots \rangle$$

obtained by iterating the operator Γ induced by the definition. In the context of the formalism used here, such an operator driven sequence is a special case of our notion of induction process obtained by applying at each stage *every* applicable rule. Our concept allows for the possibility that at some stage only one or a subset of the applicable rules is actually applied.

There are three reasons why we base our theory on this more fine-grained notion of induction process. First, we claim that it is a more faithful formalization of the way humans actually perform the induction process. E.g., to mentally compute the transitive closure of a graph from Definition 1.1, we probably do this by individual rule applications, not by applications of the operator. Second, the more fine grained induction process is needed to capture the induction process in definitions over an induction order. There, the induction process simply cannot apply all applicable rules at each stage: only those that respect the induction order can be applied. E.g., in the case of Definition 1.2, in the initial stage $\mathfrak{A}_0 = \emptyset$, every rule instance "$I \models \neg\varphi$ if $I \not\models \varphi$" applies (since $(I, \varphi) \notin \emptyset$) but application of them should be delayed until the induction process is finished with deriving $I \models \varphi$. In such cases, an operator-based induction process that applies all applicable rules at each stage, does not match with informal inductions and does not construct the defined set. Third, the more fine-grained formalization of the induction process exposes several fundamental aspects of the studied class of inductive definitions that, to the best of our knowledge, did not surface in earlier studies.

Perhaps the most striking aspect is the *non-determinism* of the induction process: by applying rules in different sequences, many induction processes can be built, even in the presence of an induction order. Now an all-important issue emerges: do all induction processes converge to the same set? If not, the definition would be ambiguous! The confluence of different induction processes of informal definitions is a fact that most of us probably take for granted; however, it is a fundamental and non-trivial property of induction. This is one of the topics that will be analyzed in this paper.

Also other aspects emerge from our study. E.g., how the induction order constrains the induction process, what the link is between the induction order and the rules, or the surprising fact that the set defined by a definition over an induction order does *not* depend on that order.

In the next section, we discuss the scope of the work of this paper, its limitations and contributions and we situate it in the broader context of mathematical and computational logic and of knowledge representation.

2 Scope of the study and related work

Related work

(Informal) definitions, including inductive ones, are fundamental in building mathematics and consequently, they have been a prime topic of research in the field of metamathematics (the mathematical study of mathematical methods). The logical study of monotone induction was started by Post (1943) and was continued in many later studies (Spector, 1961; Moschovakis, 1974a; Aczel, 1977). The study of iterated induction (which generalizes monotone induction and induction over a well-founded order) was

started by Kreisel (1963) and extended in later studies of so-called Iterated Inductive Definitions (IID) by Feferman (1970), Martin-Löf (1971), and Buchholz et al. (1981). Common to all these studies is that they focus on formal expressivity results, formal accounts of what classes of objects can be defined. In the words of Hallnäs (1991), these studies were primarily concerned with *inductive definability*, more than with *inductive definitions*.

The perspective of this paper is different. The focus is on the semantical properties of sorts of definitions that have not been fully analysed yet from a semantical point of view: definitions over an induction order and iterated inductive definitions. Our study departs from earlier studies by using a different formalization of the induction process. As a consequence, novel aspects of informal definitions emerge that were not formally studied before.

A formal "empirical" scientific study

Informal definitions exist. They appear in mathematical and scientific text. They are written, read, broadly understood, reasoned upon and computed with. Definitions are the "reality" that we study here. They are not tangible, physical objects. They are of cognitive nature and we can express and "sense" them only via language. Nevertheless, they are of mathematical, objective precision and this makes them suitable for formal scientific research.

A definition is not a physical reality, and as such some will argue that this study cannot be called an "empirical" scientific study. Nevertheless, our study shares many properties with empirical formal science. Most importantly, the theory is *falsifiable*. Any well structured informal definition of the studied class (e.g., Definitions 1.1 and 1.2) presents a potential experiment. To "execute" the experiment, one needs to establish a correspondence between the mathematical objects involved in the informal definition (parameters and potential defined sets) and structures of the vocabulary of the formal definition; then one needs to verify that the formal definition *faithfully expresses* an informal definition, in the sense defined in the introduction, then compare the mathematical object defined by the informal definition with the one defined by the formal definition. If a difference is found, the experiment refutes the theory; otherwise it confirms (corroborates) it.

In empirical sciences, there is a fundamental asymmetry between proving and disproving a theory. No number of successful experiments suffices to prove an empirical theory; but one failed experiment suffices to disprove it (Popper, 1959). Likewise, we cannot "prove" that our mathematical theory of informal definitions is correct. After all, there is no mathematical definition of what is an informal definition. Nevertheless, we are confident of our theory. In the first place, the semantic principle of the informal definitions under investigation is a solid intuition, and easy to formalize. While this is not a proof of our theory, it certainly is a compelling argument in favour for it. In the second place, contrary to physical science, any "experiment" in a study like this one is a *mathematical problem*. The correspondence between the mathematical objects of the informal definition and the structures of the formal definition, the question whether the formal definition faithfully expresses an informal definition and the correspondence between the informally defined object and the formally defined structure: they can be

analysed with mathematical methods and precision.

As such, it might be easier for us to convince the reader of the correctness of our theory than it is for a physicist to argue the correctness of a set of mathematical postulates about some physical reality. The cognitive nature of the studied objects does not prevent a precise mathematical approach to it.

Limits of the scope of the study

The scope of our study is limited in several respects. We sum up and discuss some limitations.

Many definitions in scientific and mathematical text define *(partial) functions*. The standard inductive definitions of the Fibonacci numbers and factorials, and of the truth evaluation function of propositional and predicate logic are examples. Most inductive definitions of functions are definitions over a well-founded order. Hallnäs (1991) defines and investigates a logic of inductive definitions of *(partial) functions*. In contrast, the logic that we define in this paper is to define sets.

Sets are Boolean functions and vice versa, functions are particular sets. It would not be difficult to extend our definition logic to define functions. But for now, to express definitions of n-ary (partial) functions F/n, they need to be translated to definitions of their $n+1$-ary graphs. Such transformations are routine in logic and mathematics. We see no essential difference between a definition of a function and the definition of its graph. As such, we believe that our study covers function definitions. Later this section, a small example is worked out.

The definition formalism introduced here is built on first order logic (FO): it serves to define FO predicates and the definiens of a definitional rule is a FO-formula. This is for simplicity only. In fact, all concepts in Section 3, 4, and 5 and many of them later are defined semantically in terms of the satisfaction relation of FO, and their definitions readily extend to extensions of FO equipped with a satisfaction relation (e.g., to higher-order logic, aggregate expressions, ...). Later in this section, such an example is worked out.

As a third limitation, our study covers non-inductive definitions, monotone definitions, definitions over an induction order and iterated inductive definitions. These are likely the most frequent types of definitions found in mathematical text and knowledge representation but there are other types of definitions. Not all informal definitions use the inductive constructions that we study here. Some definitions might define mathematical objects by a specific construction process expressed in the definition. Some types of definitions use a different type of induction process. Examples are inflationary induction (Moschovakis, 1974b; Gurevich and Shelah, 1985) (discussed in Section 3) and nested induction/coinduction (Sangiorgi, 2009). The latter principle is implemented in logics with nested least fixpoints such as FO(LFP) (Gurevich and Shelah, 1985), μ-calculus (Kozen, 1983) and fixpoint logics with nested least and greatest fixpoints (Bradfield, 1996; Hou, 2010).

Definitional paradoxes

One failed experiment suffices to refute a formal scientific theory but a refuted theory is not necessarily useless. E.g., Newtonian physics was soundly refuted, but it is still by far the most used physics theory. Many useful formal scientific theories are known to be only approximations of reality and to fail on borderline cases. This is not necessarily an argument to reject a theory, but rather a challenge to develop an understanding of where the theory is sufficiently precise and where are the borderline cases where it becomes unreliable.

Such phenomena will arise also in our theory. In particular, not every informal rule set constitutes a sensible informal definition. Some certainly do while others certainly don't. Where there is white and black, there is often also grey, and there is a grey zone between sensible and insensible definitions. We will argue that some quite famous "definitions" that emerged in philosophy belong to this grey or black zone: the definitional paradoxes.[2]

The informal semantics of connectives

The definition logic defined here is built on classical logic FO. Connectives and quantifiers in formal rule bodies are those of FO and they retain the informal interpretation they have in FO:

- The methodology of expressing informal rules is based on the standard interpretation of the FO connectives and quantifiers.
- The evaluation of rules during the induction process is based on standard FO semantics.

Specifically, the meaning of the negation connective in definitional rules is standard objective negation, like negation in FO.

The only truly non-standard connective in the formalism is the rule operator $\leftarrow$. A rule describes a step in the induction process, a step that produces a defined fact. As such, rules of a definition were called *productions* by Martin-Löf (1971). Rules are not truth functional; it does not even make sense to consider them as propositions (that can be true or false).

More examples

Below, a few additional informal definitions (experiments) are specified that belong to the class of definitions that we study here.

The linguistic style of expressing informal definitions as sets of cases also applies to non-recursive definitions. The definition below defines the symmetric closure SG of

[2] One aspect in which a formal study of inductive definitions is different than an empirical science is in the analysis of border cases of the system. In physics, Einstein discovered black holes by extrapolating relativity theory to its border cases; it turned out that these objects that had never been observed actually existed. But, in a formal study of a cognitive "reality" like ours, extrapolating the formalism to cases beyond what is found in scientific and mathematical text, ends up in the void: there is nothing there.

a graph G with two cases:

$$\left\{ \begin{array}{l} \forall x \forall y (SG(x,y) \leftarrow G(x,y)) \\ \forall x \forall y (SG(x,y) \leftarrow G(y,x)) \end{array} \right\}$$

A second example is in the context of the stock market. A company A *controls* a company B if the sum of the shares in B possessed by A or by any other company controlled by A, is more than 50%. To express this inductive definition, we use the sum aggregate over set expressions.

$$\left\{ \ \forall a \forall b (Contr(a,b) \leftarrow Sum(\{(x,c) \mid Shares(c,b,x) \wedge (c = a \vee Contr(a,c))\}) > 0.5) \ \right\}$$

Here, $Contr(a,b)$ has the obvious meaning that a controls b; $Shares(c,b,x)$ means that c has x shares in b. The value of a sum term $Sum(\{(x,y) \mid \varphi\})$ in structure $\mathfrak{A}$ is defined as

$$\sum_{\{(d,d') | \mathfrak{A}[x:d;y:d'] \models \varphi\}} d$$

The example definition is a monotone definition according to Definition 3.16 introduced later. It faithfully expresses the informal definition.

For a next example, the context is a (finite) transition structure $\langle S, \rightarrow \rangle$ with set of states S and transition graph $\rightarrow$. We call a state $s \in S$ *terminating* if no infinite path $s \rightarrow s_1 \rightarrow s_2 \rightarrow \ldots$ in the transition graph exists. This concept can be defined through the following inductive definition.

Definition 2.1. A state $s \in S$ is *terminating* if for each transition $s \rightarrow y$, y is terminating.

The base case of this definition is any state x without outgoing edge. This monotone inductive definition involves a universal quantifier in the definiens. Using symbol $G/2$ to express the graph and $Term/1$ to express the set of terminating states, the faithful translation of the definition in our formalism is:

$$\left\{ \ \forall x (Term(x) \leftarrow \forall y (G(x,y) \Rightarrow Term(y))) \ \right\}$$

As a second example, we define the *rank* of a terminating state s as the length of the longest path $s \rightarrow s_1 \rightarrow \ldots \rightarrow s_n$ through the transition graph. This is a partial function defined on terminating states. It can be defined inductively, over the induction order $\prec$ that is the transitive closure $\leftarrow^*$ of the inverse relation of $\rightarrow$. On the set of terminating states, $\prec$ is a strict well-founded order relation.

Definition 2.2. We define the rank of terminating states of S by induction on $\prec$:

- The rank of a terminating state $s \in L$ is the least strict upperbound of the ranks of its successors in G.

This is an informal definition of a partial function (since the rank of non-terminating states is not defined). We define the graph of the rank function in a typed variant of the definition logic:

$$\left\{ \begin{array}{l} \forall x \forall r (Rank(x,r) \leftarrow Term(x) \wedge \\ \qquad \forall y \forall r_1 (G(x,y) \wedge Rank(y,r_1) \Rightarrow r > r_1) \wedge \\ \qquad \forall r_2 (\forall y \forall r_1 (G(x,y) \wedge Rank(y,r_1) \Rightarrow r_2 > r_1) \Rightarrow r \leq r_2)) \end{array} \right\}$$

The definiens has three conditions: that x is terminating, that r is strictly larger than the rank r_1 of any successor y of x, and thirdly, that if r_2 is strictly larger than the rank of any successor of x, then r is smaller than r_2. Hence, this definition faithfully expresses the informal definition.

This is a non-monotone definition, due to the negative occurrence of $Rank$ in the middle condition. The induction process builds up $Rank$ starting from the minimal elements of $\prec$ (states without outgoing edges) for which the rank is defined to be 0 (the smallest number), and then gradually for other terminal ranks further away from these base cases. The role of the induction order is essential to obtain a correct induction proces. E.g., in the initial state of the induction process, when $Rank$ is still empty, the three conditions of the rule hold for the variable assignment $\{x = s, r = 0\}$ with s a terminating state. So untimely application of such a rule instance would derive $Rank(s) = 0$ for any terminating state s.

In the extension of FO with a minimum aggregate, the definition can be stated equivalently as:

$$\left\{ \begin{array}{l} \forall x \forall r(Rank(x,r) \leftarrow Term(x) \wedge \\ \qquad r = Minimum(\{r_1 + 1 \mid \exists y(G(x,y) \wedge Rank(y,r_1))\})) \end{array} \right\}$$

In a suitable extension of the definition logic, the body of this rule would be equivalent with the FO body above and hence, in the theory defined below, these two formal definitions of $Rank$ would have the same semantical properties: the same induction order, the same induction processes and they define the same set.

Later in this paper, one more example of an informal definition will be given, an iterated inductive definition with mixed monotone and ordered induction.

3 Formal definitions and natural inductions

Preliminaries

We introduce the concepts and notations of syntax and semantics of first order logic.

We assume an infinite supply of symbols. A symbol is either an object symbol, a predicate symbol or a function symbol. Predicate and function symbols have a unique arity n, denoting the number of arguments. Object symbols correspond to function symbols of arity 0.

The *logical symbols* are $\mathbf{t}$ (true), $\mathbf{f}$ (false), the binary equality predicate $=$, connectives $\wedge, \vee, \neg$ and the quantifiers $\exists, \forall$. Other symbols are called non-logical.

A *vocabulary* Σ is a set of non-logical symbols.

A *term* is built as usual: an object symbol is a term; if $t_1, \ldots, t_n$ are terms and f an n-ary function, then $f(t_1, \ldots, t_n)$ is a term. *Atomic formulas* or *atoms* are expressions of the form $P(t_1, \ldots, t_n)$ with P an n-ary predicate symbol and $t_1, \ldots, t_n$ terms. *Formulas* are built from atomic formulas and the logical symbols in the usual way: atoms are formulas, and if x is an object symbol and φ, ψ are formulas, then $\neg\varphi, \varphi \wedge \psi, \varphi \vee \psi, \forall x\, \varphi, \exists x\, \varphi$ are formulas. Terms and formulas are called *expressions*.

An occurrence of a symbol τ in an expression φ is *free* if τ does not occur in a subformula $\exists \tau \psi$ or $\forall \tau \psi$ of φ. The set $free(\varphi)$ is the set of all symbols that have a free

occurrence in φ. A *sentence* over Σ is a formula φ with $free(\varphi) \subseteq \Sigma$.

An occurrence of a subformula φ in ψ is called *positive* if it occurs in the scope of an even number of negations, otherwise it is called *negative*. A formula φ is *positive* with respect to a set σ of predicate symbols if there are no atoms $P(\bar{t})$ with $P \in \sigma$ that have a negative occurrence in φ.

Definition 3.1. A *structure* $\mathfrak{A}$ of vocabulary Σ consist of a non-empty set $D^{\mathfrak{A}}$ called the *domain* of $\mathfrak{A}$ and an assignment of a value $\tau^{\mathfrak{A}}$ to symbols $\tau \in \Sigma$, called the *interpretation* of τ in $\mathfrak{A}$. The value $\tau^{\mathfrak{A}}$ is an element of $D^{\mathfrak{A}}$ if τ is an object symbol, an n-ary relation over $D^{\mathfrak{A}}$ if τ is an n-ary predicate symbol, and an n-ary function on $D^{\mathfrak{A}}$ if τ is an n-ary function symbol. The interpretation of $=$ is the identity relation on $D^{\mathfrak{A}}$.

A structure $\mathfrak{A}$ of Σ is often called a Σ-structure.

The restriction of a Σ-structure $\mathfrak{A}$ to some vocabulary $\Sigma' \subseteq \Sigma$ is denoted $\mathfrak{A}|_{\Sigma'}$.

While the value $P^{\mathfrak{A}}$ of a predicate symbol is a set of n-tuples, we sometimes use it as a Boolean function, more specifically as its characteristic function. For an n-tuple $\bar{a}$ of domain elements, $P^{\mathfrak{A}}(\bar{a})$ denotes $\mathbf{t}$ if $\bar{a} \in P^{\mathfrak{A}}$, and $\mathbf{f}$ otherwise.

Definition 3.2. We define the truth evaluation function of FO by structural induction, extending $\mathfrak{A}$ to arbitrary terms and sentences over Σ using the standard inductive rules:

- $f(t_1, \ldots, t_n)^{\mathfrak{A}} = f^{\mathfrak{A}}({t_1}^{\mathfrak{A}}, \ldots, {t_n}^{\mathfrak{A}})$;
- $P(t_1, \ldots, t_n)^{\mathfrak{A}} = P^{\mathfrak{A}}({t_1}^{\mathfrak{A}}, \ldots, {t_n}^{\mathfrak{A}})$;
- $(\neg\varphi)^{\mathfrak{A}} = \neg(\varphi^{\mathfrak{A}})$,
- $(\varphi \wedge \psi)^{\mathfrak{A}} = \varphi^{\mathfrak{A}} \wedge \psi^{\mathfrak{A}}$,
- $(\varphi \vee \psi)^{\mathfrak{A}} = \varphi^{\mathfrak{A}} \vee \psi^{\mathfrak{A}}$, with $\neg$, $\wedge$, $\vee$ representing the standard boolean functions;
- $(\forall x\varphi)^{\mathfrak{A}} = minimum_{\leq_t}\{\varphi^{\mathfrak{A}[x:d]} \mid d \in D^{\mathfrak{A}}\}$, where $\mathfrak{A}[x : d]$ is the structure identical to $\mathfrak{A}$ except that $x^{\mathfrak{A}} = d$ and $\leq_t$ is the *truth order* defined by $\mathbf{f} <_t \mathbf{t}$;
- $(\exists x\varphi)^{\mathfrak{A}} = maximum_{\leq_t}\{\varphi^{\mathfrak{A}[x:d]} \mid d \in D^{\mathfrak{A}}\}$.

Let D be a non-empty set. A *domain atom* of D is a pair $(P, \bar{a})$ with P/n a predicate symbol and $\bar{a} \in D^n$. Abusing notation, we write domain atoms as atoms $P(\bar{a}), Q(\bar{b})$. We use A, B, C as mathematical variables for domain atoms. A *domain literal* is a domain atom A or its negation $\neg A$. Domain literals are denoted as L, L'. Given a structure $\mathfrak{A}$ with domain D, we define $P(\bar{a})^{\mathfrak{A}} = P^{\mathfrak{A}}(\bar{a})$. For a given set σ of predicate symbols, we denote the set of domain atoms with predicates in σ by At_D^{σ}.

For two structures $\mathfrak{A}, \mathfrak{B}$ interpreting the same vocabulary, with the same domain and interpretation of all object and function symbols, we write $\mathfrak{A} \leq_{\mathrm{t}} \mathfrak{B}$ if for every domain atom A, $A^{\mathfrak{A}} \leq_{\mathrm{t}} A^{\mathfrak{B}}$.

Formalization of informal definitions: syntax and the induction process

In this section, we formalize the syntax of the logic, when rules apply, when a definition is saturated and the induction process.

Definition 3.3. A (formal) *definition* over Σ is a set of definition rules of the form

$$\forall \bar{x}\ (P(\bar{t}) \leftarrow \phi)$$

where ϕ is a FO formula and $P(\bar{t})$ is an atomic formula over Σ such that P is not the equality predicate $=$.

We call $P(\bar{t})$ the *head* of the rule or the *definiendum*, and ϕ the *body* or the *definiens*. The connective $\leftarrow$ is called *definitional implication*. It should be distinguished from material implication. Note that this formal notion of definition does not (yet) include an induction order.

A predicate appearing in the head of a rule of a definition Δ is called a *defined predicate* of Δ; all other non-logical symbols with free occurrences in Δ are called its *parameters*. The sets of defined predicates and parameters of Δ are denoted by $def(\Delta)$ and $pars(\Delta)$, respectively. Below, a domain atom $P(\bar{a})$ of a defined predicate of Δ is called a *defined domain atom*. For simplicity, we assume that every rule is of the form $\forall \bar{x}\ (P(\bar{x}) \leftarrow \phi)$ where $\bar{x}$ is a tuple of distinct variables. Other rules $\forall \bar{x}\ (P(\bar{t}) \leftarrow \phi)$ are seen as shorthands for $\forall \bar{y}\ (P(\bar{y}) \leftarrow \exists \bar{x}(\bar{y} = \bar{t} \wedge \phi))$.

An informal rule-based definition is formally represented by choosing a suitable vocabulary Σ, and translating its informal rules rule by rule with heads representing the definiendum and bodies expressing the definiens of the informal rule. The formal definitions $\Delta_{\models}$ and Δ_{TC} were obtained this way from Definition 1.2, respectively Definition 1.1.

We will define the set defined by such formal definitions by formalizing the induction process as we discussed it above: the construction process proceeds by iterated rule application (along an induction order if there is one). As such, rules in this formalism are similar in nature to and generalize *productions* as defined by Martin-Löf (1971).

To formally define the induction process, a few auxiliary concepts are needed. First, an informal definition is always evaluated in a context of specific values for the parameters. Likewise, a formal definition is always evaluated in a *context structure*, which provides values for the parameter symbols.

Definition 3.4. We call a $pars(\Delta)$-structure $\mathcal{O}$ a *context structure* of Δ.

Given a $def(\Delta)$-structure $\mathfrak{A}$ and a context structure $\mathcal{O}$ with the same domain as $\mathfrak{A}$, we write $\mathcal{O} \circ \mathfrak{A}$ to denote the structure $\mathfrak{B}$ such that $\mathfrak{B}|_{pars(\Delta)} = \mathcal{O}$ and $\mathfrak{B}|_{def(\Delta)} = \mathfrak{A}$.

The following examples expose the context structures underlying the formal definitions Δ_{TC} and $\Delta_{\models}$. At the same time, we verify that these definitions faithfully express the corresponding informal definitions.

Example 3.5. The context structure $\mathcal{O}$ for defining the reachability relation of graph $\mathcal{G}$ on the set of vertices V, has domain V and $G^{\mathcal{O}} = \mathcal{G}$.

Combinations of graphs $\mathcal{G}$ and $\mathcal{R}$ on some domain V correspond one-to-one to Σ-structures with domain V, $G^{\mathfrak{A}} = \mathcal{G}$ and $R^{\mathfrak{A}} = \mathcal{R}$. Under this correspondence, we can verify that Δ_{TC} faithfully expresses Definition 1.1. E.g., for the inductive rule, we need to verify that in the context of arbitrary graphs $\mathcal{G}$ and $\mathcal{R}$ on domain V and corresponding $\mathfrak{A}$, that for every pair $d, e \in V$, $(d, e) \in \mathcal{R}$ if and only if $\mathfrak{A}[x : d][y : e] \models R(x, y)$ and that there exists a vertex f such that $(d, f), (f, e) \in \mathcal{R}$ if and only if

$$\mathfrak{A}[x : d][y : e] \models \exists z(R(x, z) \wedge R(z, y)).$$

This follows straight from the formal definition of FO's satisfaction relation. A similar argument holds for the base rule. Thus, Δ_{TC} faithfully expresses Definition 1.1. ▲

Example 3.6. The definition $\Delta_{\models}$ of Sat formalizes the informal Definition 1.2 of satisfaction. Its parameter symbols are $Atom, In, And, Or$ and Not. We view $\Delta_{\models}$ as a many-sorted definition, with one sort for structures and another for formulas. For any propositional vocabulary ξ, we define $PropF(\xi)$ as the set of propositional formulas over ξ and $Struct(\xi)$ as the set of propositional ξ-structures. A propositional vocabulary ξ induces the context structure $\mathcal{O}$ which is the sorted $pars(\Delta)$-structure defined as follows:

- $D^{\mathcal{O}}$ consists of two sort domains $PropF(\xi)$ and $Struct(\xi)$.
- $And^{\mathcal{O}}$ is the function that maps pairs of formulas (ψ, ϕ) to the formula $\psi \wedge \phi$. The functions $Or^{\mathcal{O}}$ and $Not^{\mathcal{O}}$ are defined in a similar vein.
- Finally, $In^{\mathcal{O}}$ is $\{(I, p) \mid I \in Struct(\xi), p \in I\}$.

As an example, two defined domain atoms for $\xi = \{P\}$ are $Sat(\{P\}, P)$ and $Sat(\{\}, P \wedge \neg P)$. The value $t^{\mathcal{O}}$ of the term $t = And(P, Not(P))$ is the formula $P \wedge \neg P$.

Given $\mathcal{O}$, there is an obvious correspondence between binary relations between structures and formulas of ξ and $def(\Delta_{\models})$-structures. Given this correspondence, it is a simple exercise to prove that the formal rules of $\Delta_{\models}$ are faithful formalisations of the corresponding informal rules. Hence, $\Delta_{\models}$ faithfully expresses Definition 1.2.

▲

From here till the end of this section, we assume the presence of a definition Δ and a context structure $\mathcal{O}$ for Δ with domain D. In the sequel, we frequently evaluate formulas with respect to structures $\mathcal{O} \circ \mathfrak{A}$. Because $\mathcal{O}$ is given and fixed, we take the liberty to write only the "variable" part and write, e.g., $\mathfrak{A} \models \varphi$ instead of $\mathcal{O} \circ \mathfrak{A} \models \varphi$, or $A^{\mathfrak{A}}$ instead of $A^{\mathcal{O} \circ \mathfrak{A}}$, etcetera.

The next definitions are formalizations of the concept of an element being *derivable* from a definition, and a set being *closed* or *saturated* under a definition.

Definition 3.7. We say that a defined domain atom $P(\bar{a})$ is *derivable* by rule $\forall \bar{x}(P(\bar{x}) \leftarrow \varphi)$ from $\mathfrak{A}$ if $\varphi^{\mathfrak{A}[\bar{x}:\bar{a}]} = \mathbf{t}$.

We say that $P(\bar{a})$ is *derivable* from $\mathfrak{A}$ (by Δ) if it is derivable by a rule of Δ from $\mathfrak{A}$. Below, we denote this by $\mathfrak{A} \vdash_{\Delta} P(\bar{a})$.

Definition 3.8. We say that $\mathfrak{A}$ is *closed* (or *saturated*) on a set S of defined domain atoms (under Δ in $\mathcal{O}$) if for every $A \in S$, $\mathfrak{A} \vdash_\Delta A$ implies $A^{\mathfrak{A}} = \mathbf{t}$. We say that $\mathfrak{A}$ is *closed* (or *saturated*) (under Δ in $\mathcal{O}$) if it is closed on $At_D^{def(\Delta)}$ (under Δ in $\mathcal{O}$).

We observe that the set of $def(\Delta)$-structures with domain $D^{\mathcal{O}}$ is isomorphic with the powerset of $At_D^{def(\Delta)}$, where the isomorphism maps such a structure $\mathfrak{A}$ to the set $\{A \in At_D^{def(\Delta)} \mid A^{\mathfrak{A}} = \mathbf{t}\}$. Under this isomorphism, the $\leq_{\mathrm{t}}$-least structure corresponds to the empty set, and the truth order $\leq_{\mathrm{t}}$ on structures corresponds to the subset relation $\subseteq$. We find it convenient to exploit this isomorphism to apply standard set theoretic operations on $def(\Delta)$-structures; e.g., denoting structures as sets of defined domain atoms, writing $A \in \mathfrak{A}$ instead of $A^{\mathfrak{A}} = \mathbf{t}$, or $\mathfrak{A} \setminus \mathfrak{A}'$ to denote the structure that interprets each predicate symbol $P \in def(\Delta)$ as $P^{\mathfrak{A}} \setminus P^{\mathfrak{A}'}$.

The following examples are the simplest sensible non-monotone definitions that we know of. They will be used as running examples through this text.

Example 3.9. Let us take the instance of Definition 1.2 of propositional satisfaction obtained by fixing $\xi = \{P\}$, by fixing the structure I to be $\{P\}$, and by limiting the formulas to those that use only the negation symbol. That is, the formulas are $P, \neg P, \neg\neg P, \ldots, (\neg)^n P, \ldots$. This instantiates the definition to

- $\{P\} \models P$
- $\{P\} \models \neg\varphi$ if $\{P\} \not\models \varphi$.

Obviously, the formulas defined to be true here are of the form $(\neg)^{2n} P$ with an even number of negations. ▲

The above definition is further simplified by transposing it to the natural numbers.

Definition 3.10. The set of even numbers is defined by induction on the standard order of natural numbers:

- 0 is even;
- $n{+}1$ is even if n is not even.

Note the correspondence with Example 3.9. This is not a common way of defining even numbers (there are much simpler ways) but it is a sensible way nevertheless.

Example 3.11. Definition 3.10 is faithfully expressed in the definition formalism as follows:

$$\Delta_{ev} = \left\{ \begin{array}{l} Even(0) \leftarrow \mathbf{t} \\ \forall x(Even(x+1) \leftarrow \neg Even(x)) \end{array} \right\} \tag{1}$$

The context structure denoted $\mathcal{O}_{ev}$ is the structure of the natural numbers, with the standard interpretation of $0, 1$ and $+$. ▲

In the following example, it is demonstrated that the non-constructive characterisation of the defined set of an inductive definition, as the least set closed under application of the rules, does not work for non-monotone definitions.

Example 3.12. Let us verify that the non-constructive characterisation of the defined set of a definition does not work for Δ_{ev}. Consider the following sets (for succinctness, we abbreviate $Even$ to Ev):

$$\{Ev(0), Ev(2), Ev(4), \dots\}$$
$$\{Ev(0), Ev(1), Ev(3), \dots\}$$
$$\{Ev(0), Ev(2), Ev(3), Ev(5), \dots\}$$
$$\dots$$
$$\{Ev(0), Ev(2), \dots, Ev(2n), Ev(2n{+}1), Ev(2n{+}3), \dots\}$$
$$\dots$$

Each of these sets represents a structure closed under Δ_{ev}. None has a strict subset that is closed under Δ_{ev}, hence each of them is minimal. Consequently, there is no least closed set. Thus, the defined set of this definition is not the least set closed under the rules. A similar phenomenon arises for the satisfaction Definition 1.2. ▲

Now we formalize the main concept of this paper: the induction process.

Definition 3.13. A *natural induction* $\mathcal{N}$ of Δ in $\mathcal{O}$ (with domain D) is a $\leq_{\mathrm{t}}$-increasing sequence $(\mathfrak{A}_\alpha)_{0\leq\alpha\leq\beta}$ of $def(\Delta)$-structures with domain D such that:

- $\mathfrak{A}_0$ is the empty structure $\emptyset$.
- For each successor ordinal $i+1 \leq \beta$, for each domain atom $A \in \mathfrak{A}_{i+1} \setminus \mathfrak{A}_i$, A is derivable from Δ in $\mathfrak{A}_i$ ($\mathfrak{A}_i \vdash_\Delta A$). We say that A is derived at i and define $\|A\|_\mathcal{N} := i$, the stage of A in $\mathcal{N}$.
- For each limit ordinal $\lambda \leq \beta$, $\mathfrak{A}_\lambda = \bigcup_{\alpha<\lambda} \mathfrak{A}_\alpha$.

We call β the length of $\mathcal{N}$, and denote $\mathfrak{A}_\beta$ as $lim(\mathcal{N})$.

Definition 3.14. A natural induction is called *terminal* if $\mathfrak{A}_\beta$ is closed under Δ (in $\mathcal{O}$).

Natural inductions will be denoted compactly as a sequence of the (disjoint) sets of atoms that are derived at each step. For instance,

$$\to \{A_1, \dots, A_n\} \to \{B_1, \dots, B_m\} \to \dots$$

derives the A_i's in step 1 and the B_j's in step 2. If such a set is a singleton we drop the brackets.

Example 3.15. Consider the formal definition Δ_{TC} formalizing the transitive closure Definition 1.1. Take context structure $\mathcal{O}$ such that $D^\mathcal{O} = \{a, b, c\}$, $G^\mathcal{O} = \{(a, a), (b, c), (c, b)\}$. All terminal natural inductions converge to $\{(a, a), (b, c), (c, b), (b, b), (c, c)\}$. For instance, the following are three different natural inductions that converge to this set:

$$\to T(a, a) \to T(b, c) \to T(c, b) \to T(b, b) \to T(c, c)$$
$$\to \{T(c, b), T(b, c)\} \to T(c, c) \to T(b, b) \to T(a, a)$$
$$\to \{T(a, a), T(b, c), T(c, b)\} \to \{T(c, c), T(b, b)\}$$

The third one is the most eager induction in the sense that it applies at each stage every applicable rule. Such a natural induction corresponds to the fixpoint computation of the operator associated with Δ_{TC}. This operator will be defined below. ▲

We want to link natural inductions of *monotone* definitions with the more standard operator-based formalization of the induction process.

Definition 3.16. We call Δ *monotone* in $\mathcal{O}$ if for all pairs of $def(\Delta)$-structures $\mathfrak{A} \subseteq \mathfrak{B}$, for all defined domain atoms A, if $\mathfrak{A} \vdash_\Delta A$ then $\mathfrak{B} \vdash_\Delta A$.

We now show that the concept of natural induction generalizes the existing operator-based formalizations of the induction process used, e.g., by Moschovakis (1974a) and Aczel (1977). Translated to our context, the induction process for a (monotone) definition Δ in $\mathcal{O}$ is formalized as the (possibly transfinite) least fixpoint construction $\emptyset, \Gamma(\emptyset), \Gamma^2(\emptyset), \ldots$ of the (monotone) operator Γ associated with Δ in $\mathcal{O}$.

Definition 3.17. The operator $\Gamma^{\mathcal{O}}_\Delta$ of Δ in context $\mathcal{O}$ is the operator of $def(\Delta)$-structures with domain $D^{\mathcal{O}}$ such that $\Gamma^{\mathcal{O}}_\Delta(\mathfrak{A}) = \{A \in At_D^{def(\Delta)} \mid \mathfrak{A} \vdash_\Delta A\}$.

Clearly, Δ is monotone in $\mathcal{O}$ if and only if $\Gamma^{\mathcal{O}}_\Delta$ is a monotone operator.

The least fixpoint construction of this operator is the (potentially transfinite) sequence:

$$\langle \mathfrak{A}_\alpha \rangle_{0 \leq \alpha \leq \beta}$$

where $\mathfrak{A}_0 = \emptyset, \mathfrak{A}_{\alpha+1} = \Gamma^{\mathcal{O}}_\Delta(\mathfrak{A}_\alpha)$, $\mathfrak{A}_\lambda = \cup_{\alpha<\lambda}\mathfrak{A}_\alpha$ for limit ordinals λ, and $\mathfrak{A}_\beta$ is a fixpoint of $\Gamma^{\mathcal{O}}_\Delta$.

It follows from Tarski's least fixpoint theorem that if $\Gamma^{\mathcal{O}}_\Delta$ is monotone, this sequence is monotonically increasing and converges to the least fixpoint of $\Gamma^{\mathcal{O}}_\Delta$. It is obvious as well that in this case, the least fixpoint construction is a special case of a natural induction; in particular, it is the most eager natural induction, the one in which a defined domain atom is derived as soon as it is derivable.

Corollary 3.18. *If Δ is monotone in $\mathcal{O}$, the least fixpoint construction of $\Gamma^{\mathcal{O}}_\Delta$ is a natural induction.*

An example of a monotone definition is the running example Δ_{TC} from Section 1. It is a *positive* definition, one in which every occurrence of a defined predicate in a rule body is positive. Positive definitions are monotone in every context structure.

Example 3.19. A monotone definition need not be positive. Three such definitions are $\{P \leftarrow P \vee \neg P\}$ and $\{P \leftarrow P \wedge \neg P\}$ and $\{P \leftarrow Q \vee (Q \wedge \neg P)\}$. ▲

It follows from Corollary 3.18 that the concept of a natural induction generalizes the least fixpoint construction. However, while the least fixpoint construction is a unique construction, a striking property of natural inductions is that that they are highly *nondeterministic*: rules can be applied in many orders, each order yielding a different natural induction. This matches our intuitive understanding that for an informal definition, there are in general many ways to construct the defined set. This is often advantageous, e.g., it allows us to pick the induction process that suits best our needs. However, there is a danger as well. From a practical point of view, it is *all-important* that different induction sequences converge to the same fixpoint, otherwise the definition would be ambiguous!

For a monotone informal definition such as Definition 1.1, the order of rule application is not important, because all sequences converge to the intended set, which is the

least relation that is closed under the rules. In the above framework of natural induction sequences, this can be proven formally.

Proposition 3.20. *Each terminal natural induction of a monotone definition Δ in $\mathcal{O}$ converges to the least $def(\Delta)$-structure $\mathfrak{A}$ that is closed under Δ in $\mathcal{O}$.*

This proposition is not difficult to prove but also follows from the general Theorem 3.50 below.

When the operator $\Gamma^{\mathcal{O}}_{\Delta}$ is not monotone, the least fixpoint construction of its operator may not converge. For such operators, Moschovakis (1974b) defined the *inflationary fixpoint construction* which is defined similarly except that $\mathfrak{A}_{\alpha+1} = \mathfrak{A}_\alpha \cup \Gamma^{\mathcal{O}}_{\Delta}(\mathfrak{A}_\alpha)$. Hence, once a defined domain atom is derived, it remains derived. Consequently, the inflationary construction yields a monotonically increasing sequence and eventually reaches a limit, called the *inflationary fixpoint*. For monotone operators, this construction coincides with the standard one, and the inflationary fixpoint is the least fixpoint.

Again, it is obvious that the inflationary fixpoint construction is the most eager natural induction of Δ in $\mathcal{O}$, the one that derives a defined domain atom as soon as it is derivable.

Corollary 3.21. *For every definition Δ and context structure $\mathcal{O}$, the inflationary fixpoint construction of $\Gamma^{\mathcal{O}}_{\Delta}$ is a natural induction.*

However, the convergence property does not hold for non-monotone definitions. In general, many natural inductions converge to different sets. The problem is that the body of a non-monotone rule may eventually become false, after it has already been true. Natural inductions that apply a rule during the "window" where its body holds will derive its head, whereas natural inductions that miss this window may not.

Example 3.22 (Continuation of Example 3.11)**.** Consider definition Δ_{ev} in the context structure $\mathcal{O}_{ev}$ of the natural numbers. The following is a natural induction:

$$\to Ev(1) \to Ev(0) \to Ev(3) \to Ev(5) \to Ev(7) \to \dots$$

Indeed, in the first step when $\mathfrak{A}_0 = \emptyset$, all instances of the rule for $Ev(x+1)$ are applicable. Here, we use it to derive $Ev(1)$. The next step applies the base rule to derive $Ev(0)$, which falsifies the condition of the rule that was applied in the first step. Next, we derive $Ev(3), Ev(5), \dots$. The limit of this natural induction is one of the unintended minimal closed sets from Example 3.12. The inflationary fixpoint construction is the terminal natural induction that converges in one step and derives evenness of all numbers:

$$\to \{Ev(n) \mid n \in \mathbb{N}\}$$

▲

Example 3.23 (Continuation of Example 3.6)**.** Consider the informal Definition 1.2 and its formalization $\Delta_{\models}$ in the context structure of the structure $\mathcal{O}$ for the singleton vocabulary $\xi = \{P\}$. There are only two structures for the vocabulary ξ, namely, $\emptyset$ and $\{P\}$. Below is an initial segment of a natural induction that derives an erroneous fact.

$$\to Sat(\{P\}, \neg P) \to Sat(\{P\}, P) \to \dots$$

In the first step, with $\mathfrak{A}_0 = \emptyset$, all instances of the rule for negation are applicable. Here, we use it to derive $Sat(\{P\}, \neg P)$. However, the next step applies the base rule to derive $Sat(\{P\}, P)$, thus falsifying the condition of the rule that was applied in the first step.

Likewise, the first step of the inflationary fixpoint construction derives all domain atoms $Sat(I, \neg\varphi)$, many of which are erroneous. This natural induction violates the induction order of Definition 1.2 and is not one of its intended induction processes. ▲

The above discussion illuminates what, in our opinion, is the essential role of the induction order in informal definitions. In definitions by induction over an induction order, the induction order serves to constrain the induction processes to ensure convergence. It does so by *delaying the application of rules until it is safe to do so*, that is, until later rule applications can no longer falsify the premise of a rule that has been applied before.

Formalization of definitions by induction over a well-founded order

We now formally define the notion of definition by induction over a well-founded order and its natural inductions.

In particular, in a context $\mathcal{O}$, we are interested in pairs $(\Delta, \prec)$ with Δ a definition and $\prec$ a strict well-founded order on $At_D^{def(\Delta)}$, referred to as the *induction order*. Recall that a strict order is irreflexive, transitive and asymmetric. A strict order $\prec$ is well-founded if it has no infinite descending chains $x_0 \succ x_1 \succ x_2 \succ \ldots$.

The following example illustrates how to formalize the induction order of an informal definition.

Example 3.24. The induction order of Definition 3.10 of even numbers is the standard order on the natural numbers. Its formalization is the order $\{Ev(n) \prec Ev(m) \mid n < m\}$. We denote it as $\prec_{ev}$. ▲

Example 3.25. Let us consider informal Definition 1.2 and its formalization, the formal definition $\Delta_{\models}$ in context structure $\mathcal{O}$ for a selected propositional vocabulary ξ.

The induction order of informal Definition 1.2 is the subformula order. The first formalization of this order that comes to mind is the strict well-founded order $\prec$ on domain atoms defined by $Sat(I, \psi) \prec Sat(J, \phi)$ if ψ is a strict subformula of ϕ. According to this order, to derive satisfaction of a formula in J, one first needs to determine the satisfaction of its subformulas in *each and every* structure I. Clearly, it suffices to determine their satisfaction in the structure J. This effect can be obtained by refining the induction order such that $Sat(I, \psi) \prec Sat(J, \phi)$ if $I = J$ and ψ is a strict subformula of ϕ. We call this the formal subformula order and denote it as $\prec_{\models}$. ▲

The induction order provided with an informal definition serves to constrain the order of rule application in natural inductions. How does this work? Intuition says that no rule should be applied to derive a fact as long as there are derivable but not yet derived facts that are strictly smaller in the induction order. For instance, assume that at some point in the induction process $I \models \varphi$ is derivable. We are allowed to make this derivation only if there is no strict subformula ψ of φ for which $I \models \psi$ is derivable but was not yet derived.

Thus, an atom $P(\bar{a})$ can be derived in the current set $\mathfrak{A}_i$ only if all strictly smaller derivable atoms in the induction order have been derived. Formally, if $\mathfrak{A}_i$ is saturated on the set $\{B \mid B \prec A\}$. This is expressed in the following definition. (Recall that the stage $\|A\|_{\mathcal{N}}$ of A in a natural induction $\mathcal{N}$ is the ordinal i such that $A \in \mathfrak{A}_{i+1} \setminus \mathfrak{A}_i$.)

Definition 3.26. A natural induction $\mathcal{N}$ *respects* $\prec$ (w.r.t. Δ and $\mathcal{O}$) if for any domain atom A derived at stage i, every atom $B \prec A$ that is derivable from $\mathfrak{A}_i$ is true in $\mathfrak{A}_i$; equivalently, if $\mathfrak{A}_i$ is saturated on $\{B \mid B \prec A\}$ (under Δ in $\mathcal{O}$).

We would expect that a natural induction that respects $\prec$ also derives atoms in this order.

Definition 3.27. We say that $\mathcal{N}$ *follows* $\prec$ if for every A and B derived by $\mathcal{N}$, $A \prec B$ implies $\|A\|_{\mathcal{N}} < \|B\|_{\mathcal{N}}$.

Example 3.28. The natural induction of Example 3.23:

$$\rightarrow Sat(\{P\}, \neg P) \rightarrow Sat(\{P\}, P) \rightarrow \dots$$

does not *respect* the formal subformula order $\prec_{\models}$. The atom $Sat(\{P\}, \neg P)$ is derived in the first step, but $Sat(\{P\}, P)$ is lower in the induction order and is derivable from $\mathfrak{A}_0$. Thus, the empty set $\mathfrak{A}_0$ is not saturated in $\{A \mid A \prec_{\models} Sat(\{P\}, \neg P)\}$. Also, this natural induction does not *follow* the formal subformula order since $Sat(\{P\}, \neg P)$ is derived before $Sat(\{P\}, P)$. ▲

Example 3.29. The natural induction of Example 3.22:

$$\rightarrow Ev(1) \rightarrow Ev(0) \rightarrow Ev(3) \rightarrow Ev(5) \rightarrow \dots$$

does not respect the formal induction order $\prec_{ev}$ and formalizes an induction process that does not respect the induction order of the informal definition formalized by Δ_{ev}. In the first step, the empty set is not saturated on $\{A \mid A \prec Ev(1)\} = \{Ev(0)\}$. ▲

In general the induction process is highly underspecified, even if an induction order is given.

Example 3.30. (Example 3.25 continued). Natural inductions of the informal Definition 1.2 will derive $I \models \varphi$ only after the satisfaction of all subformulas has been derived. This constrains the order of rule application, but much freedom is left. There are infinitely many such natural inductions. A few non-terminal ones are:

$$\rightarrow Sat(\{P\}, P) \rightarrow Sat(\{P\}, P \wedge P) \rightarrow Sat(\{P\}, \neg\neg P)$$

$$\rightarrow Sat(\{P\}, P) \rightarrow Sat(\{P\}, \neg\neg P) \rightarrow Sat(\{P\}, P \vee P)$$

Note that both natural inductions respect the subformula order and follow it. Intuition suggests that these sequences can be extended to converging terminal natural inductions, and this will be proven below. ▲

Given our experience with informal definitions, we expect some "good" properties of natural inductions that respect the induction order $\prec$: (1) that they all converge, (2) that they all *follow* the induction order, (3) that once an element is derived, it remains derivable, and (4) that in the limit, the defined set is the intended one. However, none of these properties holds right now.

The major question is related to (1). It is essential for the non-ambiguity of an informal ordered definition that all inductions that respect its induction order converge. This should be provable in our framework. However, it is straightforward to see that this is not the case. Take the empty induction order $\emptyset$ for the definition $\Delta_{\models}$ in context structure $\mathcal{O}$. This order is a strict well-founded order and all natural inductions respect it in a trivial way. As we saw in Examples 3.23 and 3.22, not all of these natural inductions converge.

As for (2), a counterexample is below.

Example 3.31. Consider the order $P \prec Q$ and definition:

$$\left\{ \begin{array}{l} Q \leftarrow \mathbf{t} \\ P \leftarrow Q \end{array} \right\}$$

Here is a terminal natural induction:

$$\rightarrow Q \rightarrow P$$

It obviously does not follow $\prec$ since $P \prec Q$. However, it does respect $\prec$. In the first step, when Q is derived, the structure $\mathfrak{A}_0 = \emptyset$ is saturated on $\{A \mid A \prec Q\} = \{P\}$, since P is not derivable. In the second step, $\mathfrak{A}_1 = \{Q\}$ is trivially saturated on $\{A \mid A \prec P\} = \{\}$. ▲

A counterexample for (3) and (4) is given below.

Example 3.32. We reconsider Δ_{ev} and $\mathcal{O}$ from Example 3.11.

$$\left\{ \begin{array}{l} Even(0) \leftarrow \mathbf{t} \\ \forall x(Even(x+1) \leftarrow \neg Even(x)) \end{array} \right\} \tag{2}$$

Recall that $\prec_{ev}$ is the order induced by the standard order on the natural numbers. That is, $Ev(n) \prec_{ev} Ev(m)$ if $n < m$. This order is *total*, and consequently, there is a unique terminal natural induction that respects it:

$$\rightarrow Ev(0) \rightarrow Ev(2) \rightarrow Ev(4) \rightarrow \ldots \rightarrow Ev(2n) \rightarrow \ldots$$

This natural induction follows $\prec_{ev}$ and constructs the set of even numbers.

Now take the following non-standard induction order:

$$Ev(1) \prec Ev(0) \prec Ev(2) \prec Ev(3) \prec \ldots$$

Also this is a total strict well-founded order. The unique terminal natural induction that respects $\prec$ is:

$$\rightarrow Ev(1) \rightarrow Ev(0) \rightarrow Ev(3) \rightarrow Ev(5) \rightarrow \ldots$$

Note that $Ev(1)$ is no longer derivable after step 2. Also, this induction clearly does not construct the intended set. ▲

In non-monotone informal definitions, we impose a well-founded induction order to obtain convergence of the induction process. However, it is clear from the above examples that in selecting the induction order, great care is required. In general, imposing an unsuitable induction order w.r.t. Δ and $\mathcal{O}$ may have a number of undesired effects as just shown.

The reason for these misbehaviours can be traced back to our earlier claim: that an induction order ensures convergence by *delaying the application of rules until it is safe to do so*. It can be seen that in the above examples, the proposed order does not achieve this. E.g., the second induction order in the above example (Ex. 3.32) allowed to derive $Ev(1)$ in the first step when it was not safe to do so; indeed, the derivation of $Ev(0)$ in the next step violates the premise of the rule that was applied to derive $Ev(1)$. This is because the proposed order $\prec$ does not reflect the *dependencies* between defined facts induced by the inductive rules. For example, while $Ev(0)$ is strictly larger than $Ev(1)$ in the proposed induction order, $Ev(1)$ is defined in terms of $Ev(0)$ and hence, if the value of $Ev(0)$ changes, this may invalidate the definiens of the rule deriving $Ev(1)$.

What emerges from this discussion is what we think to be one of the *implicit conventions* of the use of informal definitions in mathematics. Although we have never seen this explicitly stated, not every well-founded order is acceptable for use as induction order of an informal inductive definition. A "good" definition over an induction order should define the elementship of an object in the defined relation in terms of presence or absence of *strictly smaller* objects in the defined relation. This induces a constraint between the inductive rules and the induction order: a "good" induction order should *match* the dependencies amongst the defined facts induced by the inductive rules. In all above misbehaved examples, this constraint was violated. Such cases are not found in mathematical text.

We now formalize the intuition that the induction order "matches" the structure of the rules of a definition and then prove that if this condition is satisfied, the four properties (1-4) are satisfied. Intuitively, the matching condition is that defined facts may only "depend" on facts that are strictly smaller in the induction order. First, we formalize this notion of "dependence".

Let $\propto$ be an binary relation on the set $At_D^{def(\Delta)}$ of defined domain atoms. We write $\mathfrak{A}|_{\propto A}$ to denote $\mathfrak{A} \cap \{B \mid B \propto A\}$, the structure obtained from $\mathfrak{A}$ by making all domain atoms $B \not\propto A$ false.

Definition 3.33. A binary relation $\propto$ on $At_D^{def(\Delta)}$ is a *dependency relation* of Δ in $\mathcal{O}$ if for all A and all $\mathfrak{A}, \mathfrak{B}$, if $\mathfrak{A}|_{\propto A} = \mathfrak{B}|_{\propto A}$ then $\mathfrak{A} \vdash_\Delta A$ iff $\mathfrak{B} \vdash_\Delta A$.

If $\propto$ is a dependency relation, then for any defined atom A, the set $\{B \mid B \propto A\}$ is (a superset of) the set of atoms on which A depends. Indeed, in any pair of structures that coincide on this set, A is derivable in both or in none. Notice that if $\propto$ is a dependency relation, then any superset of $\propto$ is one as well.

It is convenient to extend $\propto$ to all domain literals. If L is A or $\neg A$ and L' is B or $\neg B$, then we define that $L \propto L'$ if $A \propto B$.

Example 3.34. The definition $\{P \leftarrow P\}$ has a unique dependency relation, namely $P \propto P$. For both $\{P \leftarrow P \vee \neg P\}$ and $\{P \leftarrow P \wedge \neg P\}$, the empty binary relation is

a dependency relation: P does not depend on itself. Indeed, switching the truth value of P does not affect the value of either the tautology $P \vee \neg P$ or of the contradiction $P \wedge \neg P$. ▲

That an induction order $\prec$ "matches" the rules of a definition means that $\prec$ is a dependency relation.

Example 3.35. In case of definition Δ_{ev} of even numbers and the structure $\mathcal{O}$ of Example 3.11, we see that the first order

$$Ev(0) \prec Ev(1) \prec Ev(2) \prec Ev(3) \prec \dots$$

is a dependency of Δ_{ev}, while the second order

$$Ev(1) \prec Ev(0) \prec Ev(2) \prec Ev(3) \prec \dots$$

is not. For instance, $\emptyset$ and $\{Ev(0)\}$ are identical on $\{B \mid B \prec Ev(1)\} = \emptyset$, but $\emptyset \vdash_\Delta Ev(1)$ while $\{Ev(0)\} \nvdash_\Delta Ev(1)$. ▲

Definition 3.36. We say that $\prec$ *strictly orders* Δ in $\mathcal{O}$ if $\prec$ is a strict well-founded order and a dependency relation of Δ in $\mathcal{O}$.

Example 3.37. It is an easy exercise to verify that the induction order $\prec_{\models}$ is a dependency of $\Delta_{\models}$ in the suitable context $\mathcal{O}$. Since it is a well-founded strict order, it strictly orders $\Delta_{\models}$. Likewise, $\prec_{ev}$ strictly orders Δ_{ev} in the natural numbers. ▲

Natural inductions that respect a relation $\prec$ that strictly orders Δ in $\mathcal{O}$ satisfy the good properties (1-3). (1) and (2) are shown by the following two propositions; (3) is proven later in Proposition 5.5.

Proposition 3.38. *If $\prec$ strictly orders Δ in $\mathcal{O}$ then any natural induction $\mathcal{N}$ that respects $\prec$ also follows $\prec$.*

This proposition is generalized by Proposition 6.3 and will be proven there.

Proposition 3.39. *If Δ is a definition and $\prec$ an order that strictly orders Δ in context structure $\mathcal{O}$, then terminal natural inductions that respect $\prec$ exist and all of them converge. Moreover the limit is independent of $\prec$.*

This proposition follows from the stronger Theorem 3.50.

The proposition shows that an ordered definition in which $\prec$ is a dependency of Δ unambiguously defines a set. This proposition inspires the following definitions of a definition by well-founded induction .

Definition 3.40. Let $\mathcal{O}$ be a context structure with domain D. A *definition by well-founded induction* over $\prec$ in $\mathcal{O}$ (or briefly, an *ordered definition*) is a pair $(\Delta, \prec)$ with Δ a definition and $\prec$ an order that strictly orders Δ.

Interestingly, the convergence property states that the limit is independent of the selected order. Sometimes this phenomenon can be seen in mathematical text.

Example 3.41. The Definition 1.2 defines the satisfaction relation $\models$ over the subformula order (formalized by $\prec_{\models}$) but it is not uncommon to define it over alternative induction orders. For example, we could define $\models$ by induction on the *size* of formulas. Formally, we define $Sat(I, \psi) \prec_s Sat(J, \phi))$ if $I = J$ and the size of ψ (the number of nodes in its parse tree) is strictly less than the size of ϕ. Alternatively, we may define $\models$ by induction on the *depth* of formulas, i.e., the length of the longest branch in the parse tree of ϕ. This order can be formalized similarly; let us denote its formalization as $\prec_d$. The three orders lead to three variants of Definition 1.2. Intuition suggests that they are equivalent.

It is indeed easy to verify that the formal orders $\prec_s$ and $\prec_d$ on the size and depth of formulas are supersets of $\prec_{\models}$. Hence, they are dependencies of $\Delta_{\models}$ in $\mathcal{O}$ as well. It follows from Proposition 3.39 that the natural inductions that respect them converge to the same defined set. This confirms that the three informal definitions are indeed equivalent.

This does not mean that they have the same natural inductions. For instance, reconsider the natural induction of Example 3.30:

$$\rightarrow Sat(\{P\}, P) \rightarrow Sat(\{P\}, \neg\neg P) \rightarrow Sat(\{P\}, P \vee P)$$

This one respects and follows the subformula order and the size order. However, it does not respect the depth order, since $\mathfrak{A}_1$ is not saturated on $\{B \mid B \prec Sat(\{I\}, \neg\neg P)\}$. For instance, $Sat(\{P\}, P \vee P)$ is derivable but not derived and $P \vee P$ has strictly smaller depth than $\neg\neg P$.

▲

Falsifiability In this section, we introduced and formalized two hypotheses about ordered definitions in mathematics: that their induction order is always a dependency of the definition (Definition 3.33) and how an induction process respects the induction order (Definition 3.26). The former hypothesis stems from the view that such a definition defines elements of the defined sets in terms of strictly smaller elements. Certainly, these hypotheses cannot be proven but they are falsifiable in concrete "experiments". They are satisfied in the satisfaction Definition 1.2 and the evenness Definition 3.10, in their formal representations and in all other informal definitions over an induction order that the authors are aware of. Also, as formally proven in Proposition 3.39, definitions and inductions satisfying these hypotheses possess the indispensible confluence property: all induction processes converge to the same limit. Thus, such definitions *define* a set.

Generalizing monotone and ordered definitions.

There is an obvious similarity between Propositions 3.20 and 3.39 of the confluence of natural inductions of monotone and ordered definitions. However, neither is a generalization of the other. Not all monotone definitions are ordered. For instance, for the definition Δ_{TC} of transitive closure in $\mathcal{O}$, there is no $\prec$ that strictly orders Δ_{TC} in $\mathcal{O}$. Indeed, due to the transitivity rule, all defined domain atoms depend on each other; the only dependency relation is the total one and this is not a strict order.

We now define the more general class of *iterated inductive definitions*, which encompasses all ordered definitions as well as all monotone definitions. We will then prove a theorem for this more general class that generalizes both of the earlier results.

The general idea of iterated inductive definitions is that they admit a dependency $\propto$ that is not a strict order; however, if atoms A, B depend on each other (that is, $A \propto B \propto A$), then they depend *monotonically* on each other: deriving B may switch A from underivable to derivable but not from derivable to underivable; A's effect on B is similar.

For a given dependency $\propto$, we define $A \prec_\propto B$ if $A \propto B$ and $B \not\propto A$. If $\propto$ is transitive, then $\prec_\propto$ is a strict order. In that case, $\prec_\propto$ divides the set of domain atoms into a set of strictly ordered "layers" such that, for all A, B, if $A \prec_\propto B$, then A is in a strictly lower layer than B, and if $A \propto B \propto A$, they are in the same layer. If moreover, the layers form a well-founded order then we have an iterated inductive definition.

Natural inductions of an iterated inductive definition proceed along the order $\prec_\propto$. Such a natural induction closes layer by layer using monotone "sub-inductions" that take place inside a single layer, and starts a new monotone induction in the next layer as soon as one layer is saturated. To ensure this behaviour, the same condition is imposed on natural inductions as for an ordered definition: an atom A may be derived at step i only if $\mathfrak{A}_i$ is saturated on $\{B \mid B \prec_\propto A\}$.

We now formalize these ideas.

Definition 3.42. A relation $\propto$ is a *monotone dependency relation* of Δ in $\mathcal{O}$ if for all defined A, for all $\mathfrak{A}, \mathfrak{B}$ such that $\mathfrak{A}|_{\prec_\propto A} = \mathfrak{B}|_{\prec_\propto A}$ and $\mathfrak{A}|_{\propto A} \subseteq \mathfrak{B}|_{\propto A}$, if $\mathfrak{A} \vdash_\Delta A$ then $\mathfrak{B} \vdash_\Delta A$.

Proposition 3.43. *If $\propto$ is a monotone dependency relation of Δ in $\mathcal{O}$ then $\propto$ is a dependency relation of Δ in $\mathcal{O}$.*

Proof. If $\mathfrak{A}|_{\propto A} = \mathfrak{B}|_{\propto A}$, then $\mathfrak{A} \vdash_\Delta A$ implies $\mathfrak{B} \vdash_\Delta A$ and vice versa. □

Just as for dependencies, it is easy to see that any superset of a monotone dependency is a monotone dependency as well. In particular, the transitive closure of a monotone dependency is one. Thus, any definition that admits a monotone dependency admits a transitive monotone dependency.

Definition 3.44. A relation $\propto$ *monotonically orders* Δ in $\mathcal{O}$ if $\propto$ is transitive, $\prec_\propto$ is a strict well-founded order and $\propto$ is a monotone dependency relation of Δ in $\mathcal{O}$.

Definition 3.45. We say that a natural induction $\mathcal{N}$ *respects* (*follows*) a transitive relation $\propto$ if it respects (follows) $\prec_\propto$ according to Definition 3.26.

Thus, if $\mathcal{N}$ respects $\propto$ and $A \in \mathfrak{A}_{i+1} \setminus \mathfrak{A}_i$ then $\mathfrak{A}_i|_{\prec_\propto A}$ is saturated. If $\mathcal{N}$ follows $\propto$ then for every A and B derived by $\mathcal{N}$, $A \prec_\propto B$ implies $\|A\|_\mathcal{N} < \|B\|_\mathcal{N}$.

We have already defined the concept of a monotone and ordered definition in context structure $\mathcal{O}$. Now, we also define the concept of an iterated inductive definition (in $\mathcal{O}$).

Definition 3.46. A definition Δ is a *definition by iterated induction* over $\propto$ in $\mathcal{O}$ if $\propto$ monotonically orders Δ in $\mathcal{O}$.

We first show that iterated inductive definitions generalize monotone and ordered definition. For a monotone definition, the entire set of all domain atoms can serve as a single layer. Let $\propto_t$ denote the total binary relation on $At_\Delta^{def(\Delta)}$. Note that $\prec_{\propto_t} = \emptyset$.

Proposition 3.47. *A definition Δ is monotone in $\mathcal{O}$ iff Δ is a definition by iterated induction over $\propto_t$ in $\mathcal{O}$. A natural induction of Δ in $\mathcal{O}$ (trivially) respects $\propto_t$.*

Proof. Since $\prec_{\propto_t} = \emptyset$, the condition that Δ is an iterated definition over $\propto_t$ in $\mathcal{O}$ collapses to the condition that for all A, $\mathfrak{A}, \mathfrak{B}$ such that $\mathfrak{A} \subseteq \mathfrak{B}$, if $\mathfrak{A} \vdash_\Delta A$ then $\mathfrak{B} \vdash_\Delta B$. This is precisely the monotonicity condition. □

Example 3.48. Consider the formal definition Δ_{TC} of transitive closure and the context structure $\mathcal{O}$ with domain $\{a, b, c\}$ of Example 3.15. The total binary relation of $At_D^{def(\Delta)}$ is the one and only dependency relation of Δ_{TC} in $\mathcal{O}$. ▲

Proposition 3.49. *For a binary relation $\propto$, a definition Δ is a definition by well-founded induction over $\propto$ in $\mathcal{O}$ iff Δ is by iterated induction over $\propto$ in $\mathcal{O}$ and in addition, $\propto$ is irreflexive and asymmetric (and hence, a strict order).*

Proof. Obvious from the definitions. □

Given that monotone and ordered definitions are special cases of iterated inductive definitions, the following proposition presents a generalization of both Proposition 3.20 and Proposition 3.39.

Theorem 3.50. *Assume that Δ is by iterated induction over $\propto$ in $\mathcal{O}$. Then terminal natural inductions that respect $\propto$ exist and all converge. Moreover, the limit is independent of $\propto$.*

This theorem follows from Theorem 5.7 and will be proven below.

Definition 3.51. The *structure defined by* a definition Δ by iterated induction over $\propto$ in $\mathcal{O}$ is the limit of any terminal natural induction that respects $\propto$.

Informal iterated inductive definitions Above, formal iterated inductive definitions were introduced as a mathematical generalization of monotone and ordered definition. In this section we discuss their application in mathematical text.

Quite a few definitions in mathematical text contain iterated applications of nested monotone induction. However, they are only rarely formulated as sets of informal rules. To phrase them, formal scientists typically use other tools from their toolbox, such as fixpoints of operators. A well-known iterated inductive definition is the alternating fixpoint definition of the well-founded model (Van Gelder, 1993). In this definition, the well-founded model of logic program Π is characterised as the limit of an alternating fixpoint construction of an anti-monotone operator $\mathcal{A}$. This operator, called the stable operator $\mathcal{A}$ of Π is defined on structures $\mathfrak{A}$ by defining $\mathcal{A}(\mathfrak{A})$ as the least fixpoint of some monotone operator $\lambda x T(x, \mathfrak{A})$ associated to Π (essentially, the four-valued immediate consequence operator of Π). This is an iterated induction in the sense that each of the steps in alternating sequence involves itself a monotone inductive construction.

A rare case where iterated induction is explicitly available in rule form is in the definition of a *stable theory* (Marek, 1989) which is a set of propositional modal logic formulas closed under the standard inference rules and two additional ones:

$$\frac{\vdash \psi}{\vdash K\psi} \qquad \frac{\nvdash \psi}{\vdash \neg K\psi}$$

The second is a non-monotone rule. The set is computed by iterated induction for increasing modal nesting depth of modal formulas.

In the following example, we rephrase the definition of the satisfaction relation of multi-agent modal logic as an informal iterated inductive definition.

Example 3.52. Consider the multi-agent modal logic with a finite set of agents $\mathcal{A}$, the standard propositional connectives, for each agent $a \in \mathcal{A}$ the epistemic operator K_a, and for each group of agents $g \subseteq \mathcal{A}$ the common knowledge operator C_g and its dual operator DC_g. These operators satisfy the standard condition $C_g\varphi \Leftrightarrow \neg DC_g\neg\varphi$.

The satisfaction relation is defined in terms of (multi-agent) Kripke structures and worlds. A multi-agent Kripke structure is a tuple $\mathcal{K} = \langle \mathcal{W}, \xi, L, \mathcal{A}, R\rangle$ with $\mathcal{W}$ a set of worlds, ξ a propositional vocabulary, $L : \mathcal{W} \to 2^\xi$ a function from worlds to ξ -structures, $\mathcal{A}$ the set of agents, and $R \subseteq \mathcal{W} \times \mathcal{A} \times \mathcal{W}$ the accessibility relation: if $(w_1, a, w_2) \in R$ then according to agent a, world w_2 is accessible from world w_1.

The formula $C_g\varphi$ holds in a world w if every finite path from w through the union of the accessibility relations of agents in g ends in a world w' that satisfies φ. Correspondingly, $DC_g\varphi$ holds if at least one such a path exists; that is, if a world satisfying φ is *reachable* in the combined accessibility relation of the agents of g. This reachability condition can be expressed through a monotone inductive rule. $C_g\varphi$ can be defined in terms of $DC_g\neg\varphi$ using a non-monotone rule.

Below, we specify the informal definition together with its monotone dependency relation $\propto$. The right column specifies for each inductive rule the dependencies that it generates, and whether these dependencies are cyclic ($A \propto B$) or not ($A \prec_\propto B$). In this definition, the domain atoms are of the form $\mathcal{K}, w \models \varphi$. For brevity, we drop the argument $\mathcal{K}$.

– $\mathcal{K}, w \models p$ if $p \in \xi$ and $p \in L(w)$	
– $\mathcal{K}, w \models \neg\varphi$ if $\mathcal{K}, w \not\models \varphi$	$(w \models \varphi) \prec_\propto (w \models \neg\varphi)$
– $\mathcal{K}, w \models \psi \wedge \phi$ if $\mathcal{K}, w \models \psi$ and $\mathcal{K}, w \models \phi$	$(w \models \psi) \prec_\propto (w \models \psi \wedge \phi)$, $(w \models \phi) \prec_\propto (w \models \psi \wedge \phi)$
– $\mathcal{K}, w \models \psi \vee \phi$ if $\mathcal{K}, w \models \psi$ or $\mathcal{K}, w \models \phi$ (or both)	$(w \models \psi) \prec_\propto (w \models \psi \vee \phi)$, $(w \models \phi) \prec_\propto (w \models \psi \vee \phi)$
– $\mathcal{K}, w \models K_a\varphi$ if there exists $(w, a, w') \in R$ such that $\mathcal{K}, w' \models \varphi$	$(w' \models \varphi) \prec_\propto (w \models K_a\varphi)$
– $\mathcal{K}, w \models DC_g\varphi$ if there exists $(w, a, w') \in R$ such that $a \in g$ and $\mathcal{K}, w' \models \varphi$	$(w' \models \varphi) \prec_\propto (w \models DC_g\varphi)$
– $\mathcal{K}, w \models DC_g\varphi$ if there exists $(w, a, w') \in R$ such that $a \in g$ and $\mathcal{K}, w' \models DC_g\varphi$	$(w' \models DC_g\varphi) \propto (w \models DC_g\varphi)$
– $\mathcal{K}, w \models C_g\varphi$ if $\mathcal{K}, w \not\models \neg DC_g\varphi$.	$(w \models DC_g\neg\varphi) \prec_\propto (w \models C_g\varphi)$.

The relation $\propto$ is the transitive closure of the collection of all tuples specified in the right column, for all w, w' and all formulas $\varphi, \psi, \phi, K_a\varphi, DC_g\varphi, C_g\varphi$. $\propto$ is not a strict

order since it contains cycles. The cycles are the dependencies $(w' \models DC_g\varphi) \propto (w \models DC_g\varphi)$ induced by the second, monotone rule for $DC_g\varphi$. It can be easily verified that the strict order $\prec_\propto$ is well-founded.

The definition contains non-monotone rules for $\neg\varphi$ and for $C_g\varphi$. The definition is not an ordered definition, since $\propto$ is not a strict order. However, $\propto$ is a monotone dependency relation of this iterated inductive definition. Consequently, all natural inductions that respect $\propto$ converge to the intended relation. Hence, this is a well-defined informal iterated inductive definition.

▲

Also in knowledge representation, one sometimes finds natural applications of iterated inductive definitions that can be faithfully expressed as rule sets. For instance, Denecker and Ternovska (2007) argued that dynamic systems with cyclic ramifications can be naturally described using iterated inductive definitions.

Summary: implications for informal definitions The formalization of definitions in this section exposes and proves several fundamental properties of informal definitions.

First, that the "non-constructive" characterization of the defined set as the least set satisfying the rules, is incorrect in case of non-monotone (ordered or iterated) definitions.

Second, that the induction process, seen as the iterated application of rules, is highly non-deterministic, and therefore that convergence is all-important. In mathematical practice, we typically take this property for granted. In fact, it is not trivial at all. It is a fundamentally important property of inductive definitions, of great pragmatical importance.

Third, we formalized how the induction order is to be used in the induction process in the concept of a natural induction respecting an induction order.

Fourth, in mathematical texts, we have a certain degree of freedom when it comes to choosing the induction order for an inductive definition. Nevertheless, the order is far from arbitrary and needs to match the structure of the rules. This match was formalized in the concept of *dependency*. Our exposition clarifies the role and nature of the induction order, the match with the definitional rules and how the induction order constrains the order of rule application. We were then able to state Theorem 3.50 that all natural inductions that respect such a relation converge (the proof is given in the next sections).

Last but not least, it also appears from Theorem 3.50 that the choice of the induction order is irrelevant as long as it matches the rules. The order does not affect the semantics of the definition. In view of this, one may wonder why an induction order is specified at all in mathematical text. This will be explored in the next sections.

Related work on iterated induction Iterated inductive definitions were studied in (Kreisel, 1963; Feferman, 1970; Martin-Löf, 1971; Buchholz et al., 1981). In the formalisms of (Kreisel, 1963; Feferman, 1970; Martin-Löf, 1971), a strict syntactical stratification condition on rule sets ensure that the rule set Δ admits a monotone dependency in *every context* $\mathcal{O}$ and hence, is an iterated definition in every $\mathcal{O}$. This condition is

similar to the notion of stratification in logic programming (Apt et al., 1988). A disadvantage of this approach is that many (nonmonotone) informal and formal definitions are sensible definitions in one context but not in another. E.g., the evenness Definition 3.10 and its faithful representation Δ_{ev} are sensible definitions in the context of the natural numbers, but not in the context of the integer numbers. Indeed, in the integer numbers, the only dependency of this definition is still the standard order but this order is not well-founded in the integers. Also the satisfaction Definition 1.2 is not an ordered definition in every context $\mathcal{O}$.

A more general approach is the logic of iterated induction (IID) presented by Buchholz et al. (1981). There, an iterated inductive definition is expressed via a second order logic formula that expresses a definition Δ and, independently, an induction order $\prec$. They use this logic to study proof-theoretic strength and expressivity of iterated definitions. In the IID formalism, the order can be chosen independently of the definition; there is no requirement similar to our notion of dependency. We showed that the risks of choosing an order that does not match with the definition are that (1) there is no convergence of different induction processes, and (2) that an unintended set is constructed. The first problem is avoided in the logic of Buchholz et al. (1981). Essentially, the second order formula constrains the induction process to a single process. As for the second problem, it is possible in this formalism to encode an induction order that does not match the rules. For example, one can encode the definition Δ_{ev} with the non-matching order $Ev(1) \prec Ev(0) \prec Ev(2) \prec \ldots$, in which case the unintended set $\{Ev(1), Ev(0), Ev(3), Ev(5), \ldots\}$ is constructed.

In some sense, the IID logic is more general than the formalism here, since by selecting different induction orders for the same rule set, different induction processes and different defined sets can be obtained. If our hypothesis about the link between rules and induction order is correct, this extra expressivity does not cover useful ground, moreover it poses two disadvantages. First, formally expressing an induction order in the logic might be as complex as expressing the definition itself, if not more. Second, it also makes the knowledge representation process more error-prone, if there is no way to prevent that an order is encoded that does not match with the definition. To have to express the induction order seems like a needless complication of the knowledge representation process.

To us it seems preferable to design a logic of definitions in which only the rules need to be represented and the order is left implicit. Indeed, Theorem 3.50 gives us license to do this, because it shows that all induction orders that fit the structure of the rules of Δ produce the same unique limit of their terminal natural inductions. Nevertheless, it could be useful to express an induction order as a "parity check' for the correctness of the definition.

4 Safe natural inductions

In the previous section, we argued that the role of the induction order is to delay the application of a rule until it is *safe* to do so, i.e., until later rule applications cannot violate the premise of the rule anymore. In this section, we formalize this intuition and prove its correctness.

To define safe natural inductions, we need a slightly extended notion of natural induction that starts from an arbitrary $def(\Delta)$-structure $\mathfrak{A}$ rather than from $\emptyset$.

Definition 4.1. We define a *natural induction $\mathcal{N}$ of Δ* **from a** *$def(\Delta)$*-**structure** $\mathfrak{A}$ *in* $\mathcal{O}$ in the same way as Definition 3.13 except that $\mathfrak{A}_0 = \mathfrak{A}$.

We introduce the following notations. Given a natural induction $\mathcal{N}$ with limit $lim(\mathcal{N}) = \mathfrak{A}$ and a natural induction $\mathcal{N}'$ from $\mathfrak{A}$, their composition $\mathcal{N} + \mathcal{N}'$ is obtained by appending $\mathcal{N}'$ after $\mathcal{N}$. Clearly, the result is a natural induction. Also, for a natural induction $\mathcal{N} = \langle \mathfrak{A}_\alpha \rangle_{0 \leq \alpha \leq \beta}$ and $0 \leq i \leq j \leq \beta$, we write $\mathcal{N}_{i \to j}$ to denote the segment $\langle \mathfrak{A}_i, \mathfrak{A}_{i+1}, \ldots, \mathfrak{A}_j \rangle$ of $\mathcal{N}$. This is a natural induction from $\mathfrak{A}_i$.

Definition 4.2. A defined atom A is *safely derivable* by Δ in structure $\mathfrak{A}$ if $\mathfrak{A} \vdash_\Delta A$ and for each natural induction $\mathcal{N}$ of Δ from $\mathfrak{A}$, it holds that $lim(\mathcal{N}) \vdash_\Delta A$. The set of safely derivable atoms from $\mathfrak{A}$ is denoted $Safe_\Delta(\mathfrak{A})$.

Definition 4.3. We call A *strictly underivable* by Δ in structure $\mathfrak{A}$ if for each natural induction $\mathcal{N}$ of Δ from $\mathfrak{A}$, it holds that $lim(\mathcal{N}) \nvdash_\Delta A$. The set of strictly underivable atoms from $\mathfrak{A}$ is denoted $Underivable^*_\Delta(\mathfrak{A})$.

We will see that in a terminal natural induction, every atom that is safely derivable at some stage, is eventually derived. An atom that is strictly underivable at some stage is never derived.

Definition 4.4. The structure $\mathfrak{B}$ is safely derivable from $\mathfrak{A}$ if $\mathfrak{A} \subseteq \mathfrak{B} \subseteq \mathfrak{A} \cup Safe_\Delta(\mathfrak{A})$. Equivalently, if $\mathfrak{A} \subseteq \mathfrak{B}$ and every $A \in \mathfrak{B} \setminus \mathfrak{A}$ is safely derivable in $\mathfrak{A}$.

Definition 4.5. A natural induction $\mathcal{N} = \langle \mathfrak{A}_\alpha \rangle_{\alpha \leq \beta}$ from $\mathfrak{A}$ is safe if for each $\alpha < \beta$, $\mathfrak{A}_{\alpha+1}$ is safely derivable from $\mathfrak{A}_\alpha$.

An obvious property of safe natural inductions is that any atom A that is derived at some stage i remains derivable at later stages. The following proposition states that in a natural induction the sets of safely derivable and of strictly underivable defined atoms grow monotonically.

Proposition 4.6. *If $\mathcal{N} = \langle \mathfrak{A}_\alpha \rangle_{\alpha \leq \beta}$ is a natural induction from $\mathfrak{A}$, then for all $0 \leq i < j \leq \beta$, $Safe_\Delta(\mathfrak{A}_i) \subseteq Safe_\Delta(\mathfrak{A}_j)$ and $Underivable^*_\Delta(\mathfrak{A}_i) \subseteq Underivable^*_\Delta(\mathfrak{A}_j)$.*

Proof. Assume that $A \in Safe_\Delta(\mathfrak{A}_i)$ is not safely derivable in $\mathfrak{A}_j$. Let $\mathcal{N}'$ be a natural induction from $\mathfrak{A}_j$ such that $lim(\mathcal{N}') \nvdash_\Delta A$. Then $\mathcal{N}_{i \to j} + \mathcal{N}'$ is a natural induction from $\mathfrak{A}_i$ to $lim(\mathcal{N}')$. Hence, A is not safely derivable in $\mathfrak{A}_i$. Contradiction. The case for underivability is similar. □

Proposition 4.7. *Let $\mathcal{N} = \langle \mathfrak{A}_\alpha \rangle_{\alpha \leq \beta}$, $\mathcal{N}' = \langle \mathfrak{B}_\alpha \rangle_{\alpha \leq \gamma}$ be two safe natural inductions from the same structure $\mathfrak{A}$. For every $i \leq \beta, j \leq \gamma$ it holds that if $i + 1 \leq \beta$ then $\mathfrak{A}_{i+1} \cup \mathfrak{B}_j$ is safely derivable from $\mathfrak{A}_i \cup \mathfrak{B}_j$ and if $j + 1 \leq \gamma$ then $\mathfrak{A}_i \cup \mathfrak{B}_{j+1}$ is safely derivable from $\mathfrak{A}_i \cup \mathfrak{B}_j$.*

Proof. The product order $\leq$ for ordinal pairs (given by $(i, j) \leq (k, l)$ if $i \leq k, j \leq l$) is a well-founded order, hence every set of such pairs contains minimal elements in this order.

Assume towards contradiction that pairs $(i,j) \leq (\beta,\gamma)$ exist that contradict the proposition, and let (i,j) be a minimal such pair in the product order. Hence, either $\mathfrak{A}_{i+1} \cup \mathfrak{B}_j$ exists and is not safely derivable from $\mathfrak{A}_i \cup \mathfrak{B}_j$, or $\mathfrak{A}_i \cup \mathfrak{B}_{j+1}$ exists and is not safely derivable.

Assume that it is the first case. Thus, $\mathfrak{A}_{i+1} \cup \mathfrak{B}_j$ exists (i.e., $i+1 \leq \beta$) and $\mathfrak{A}_{i+1} \cup \mathfrak{B}_j$ is not safely derivable from $\mathfrak{A}_i \cup \mathfrak{B}_j$: at least one domain atom $A \in \mathfrak{A}_{i+1} \setminus \mathfrak{A}_i$ is safely derivable from $\mathfrak{A}_i$ but not from $\mathfrak{A}_i \cup \mathfrak{B}_j$. By the minimality of (i,j), the sequence $\mathcal{N}" = \langle \mathfrak{A}_i \cup \mathfrak{B}_\alpha \rangle_{0 \leq \alpha \leq j}$ contains only safe derivations and hence, it is a natural induction from $\mathfrak{A}_i$ to $\mathfrak{A}_i \cup \mathfrak{B}_j$. Since the set of safely derivable domain atoms grows in this sequence (Proposition 4.6), A is safely derivable from $\mathfrak{A}_i \cup \mathfrak{B}_j$. We obtain the contradiction. The second case is obtained by symmetry. □

It follows that every path in the "matrix" of structures $\mathfrak{A}_i \cup \mathfrak{B}_j$ obtained by incrementing at each step either i or j by 1, is a safe natural induction.

Definition 4.8. $\mathcal{N}$ is *safe-terminal* if $\mathcal{N}$ is safe and $Safe_\Delta(lim(\mathcal{N})) \subseteq lim(\mathcal{N})$.

In other words, a safe natural induction $\mathcal{N}$ is safe-terminal if it cannot be extended to a larger safe natural induction.

Theorem 4.9. *All safe-terminal natural inductions converge to the same structure.*

Proof. Take two safe-terminal natural inductions $\mathcal{N} = \langle \mathfrak{A}_\alpha \rangle_{\alpha \leq \beta}$, $\mathcal{N}' = \langle \mathfrak{B}_\alpha \rangle_{\alpha \leq \gamma}$. Consider the sequence $\langle \mathfrak{C}_\alpha \rangle_{\alpha \leq \beta+\gamma}$ where $\mathfrak{C}_\alpha = \mathfrak{A}_\alpha$ if $\alpha \leq \beta$ and $\mathfrak{C}_{\beta+\alpha} = \mathfrak{A}_\beta \cup \mathfrak{B}_\alpha$ if $\alpha \leq \gamma$. This sequence corresponds to the path through the matrix going first from $\emptyset$ to $\mathfrak{A}_\beta$ following $\mathcal{N}$ and then from $\mathfrak{A}_\beta$ to $\mathfrak{A}_\beta \cup \mathfrak{B}_\gamma$. By Proposition 4.7, this is a safe natural induction. Since $Safe_\Delta(\mathfrak{A}_\beta) = \emptyset$, the sequence is constant starting from $\mathfrak{C}_\beta$; i.e., $\mathfrak{C}_{\beta+\alpha} = \mathfrak{A}_\beta$ for all $\alpha \leq \gamma$. Hence $\mathfrak{A}_\beta = \mathfrak{A}_\beta \cup \mathfrak{B}_\gamma$ and $\mathfrak{B}_\gamma \subseteq \mathfrak{A}_\beta$. By a symmetrical argument also the converse inclusion holds. □

Definition 4.10. The structure *safely defined* by Δ in $\mathcal{O}$ is the limit of any safe-terminal natural induction of Δ in $\mathcal{O}$.

The results of this section show that imposing safety on natural inductions ensures confluence. We still need to show that natural inductions that respect a suitable induction order are safe. This is done in the next section.

5 Existence and confluence of natural inductions in ordered and iterated definitions

We now explore basic properties of informal inductive definitions. Often they are "evident" to us; some are critical for practical reasoning on informal inductive definitions. Nevertheless, they are non-trivial and here we prove them in the context of the formal framework.

Existence of terminal natural inductions We show that the condition on $\propto$ in the definition of ordered and iterated inductive definitions that $\prec_\propto$ is a strict well-founded order, serves to ensure that a sound non-terminal natural induction can always be extended to a terminal one. Thus, a sound induction process cannot "stall" in the middle.

Proposition 5.1. *Let Δ be a definition, $\mathcal{O}$ a context structure, $\propto$ a transitive binary relation on $At_D^{def(\Delta)}$. If $\prec_\propto$ is a well-founded strict order, then any natural induction $\mathcal{N}$ that respects $\propto$ can be extended to a terminal natural induction that respects $\propto$.*

Note that for this proposition to hold, it is not necessary that $\propto$ is a dependency relation of Δ in $\mathcal{O}$ but only that $\prec_\propto$ is a strict well-founded order.

Proof. Assume towards contraction that $\mathcal{N}$ respects $\propto$ but cannot be extended to a terminal natural induction that respects $\propto$. Without loss of generality, we may assume that $\mathcal{N}$ is a maximal such a sequence, that is, it respects $\propto$ but cannot be extended to a natural induction that respects $\propto$. Let $\mathcal{N}$'s last element be $\mathfrak{A}_\beta$. Since $\mathcal{N}$ is non-terminal, there exists at least one A such that $\mathfrak{A}_\beta \vdash_\Delta A$ and $A^{\mathfrak{A}_\beta} = \mathbf{f}$. Consider the set of all such atoms. Since $\prec_\propto$ is a well-founded order, this set has at least one $\prec_\propto$-minimal element A. Due to its minimality, $\mathfrak{A}_\beta$ is saturated on $\{B \in At_D^{def(\Delta)} \mid B \prec_\propto A\}$. Hence, the extension of $\mathcal{N}$ with $\mathfrak{A}_\beta \cup \{A\}$ is a natural induction that respects $\propto$. Contradiction. □

The condition of well-foundedness of $\prec_\propto$ is necessary. E.g., when interpreting the definition Δ_{ev} in the context of the integer numbers instead of the natural numbers, the strict order $\{Ev(n) \prec Ev(m) \mid n < m \in \mathbb{Z}\}$ is a dependency of the definition Δ_{ev}. Nevertheless, the definition does not have non-trivial natural inductions that respect this order, and this is due to the fact that the order is not well-founded.

Confluence of terminal natural inductions.

Proposition 5.2. *A terminal safe natural induction is safe-terminal.*

Proof. Trivial since safely derivable atoms are derivable. □

All safe-terminal natural inductions converge. Safe natural inductions that are terminal are safe-terminal. Thus, to prove the confluence of terminal natural inductions of Δ respecting a suitable $\propto$, it suffices to prove that natural inductions respecting $\propto$ are safe.

Let $\propto$ be an arbitrary binary relation on the defined domain atoms.

Definition 5.3. A set S of domain atoms is $\propto$-*closed* if for all $A \in S$, for all $B \propto A$, it holds that $B \in S$.

We observe that if $\propto$ is transitive then for every A, $\{B \mid B \propto A\}$ is $\propto$-closed.

The next proposition states that once some intermediate structure $\mathfrak{A}_i$ in a natural induction $\mathcal{N}$ is saturated on a $\propto$-closed set S, then the value and derivability of atoms of S does not change anymore later in $\mathcal{N}$.

Proposition 5.4. *Assume $\propto$ monotically orders Δ in $\mathcal{O}$. Let $\mathfrak{A}$ be a structure and $\mathcal{N}$ an arbitrary natural induction from $\mathfrak{A}$.*

1. *If $\mathfrak{A}$ is saturated on a $\propto$-closed set S, then $lim(\mathcal{N}) \cap S = \mathfrak{A} \cap S$.*

2. *If $\mathfrak{A}$ is saturated on $\{B \mid B \propto A\}$ then if $\mathfrak{A} \vdash_\Delta A$ then $lim(\mathcal{N}) \vdash_\Delta A$.*

Proof. Let $\mathcal{N}$ be of the form $\langle \mathfrak{A}_\alpha \rangle_{0 \leq \alpha \leq \beta}$ with $\mathfrak{A}_0 = \mathfrak{A}$.

(1) We prove $\mathfrak{A}_\alpha \cap S = \mathfrak{A} \cap S$ for all $\alpha \leq \beta$ by induction on α.

It trivially holds for $\alpha = 0$ since $\mathfrak{A}_0 = \mathfrak{A}$.

Assume it holds for α. For any $A \in S$, it holds that $\{B \mid B \propto A\} \subseteq S$, hence $\mathfrak{A}_\alpha \cap \{B \mid B \propto A\} = \mathfrak{A} \cap \{B \mid B \propto A\}$. Since $\propto$ is dependency of Δ, it follows that $\mathfrak{A}_\alpha \vdash_\Delta A$ iff $\mathfrak{A} \vdash_\Delta A$. Since $\mathfrak{A}$ is saturated on S, it is saturated on $\{B \mid B \propto A\}$. Hence, if $\mathfrak{A}_\alpha \vdash_\Delta A$ then $A \in \mathfrak{A}$. Consequently, $\mathfrak{A}_{\alpha+1} \cap S = \mathfrak{A} \cap S$.

Finally, assume that for limit ordinal λ, for every $\alpha < \lambda$, it holds that $\mathfrak{A}_\alpha \cap S = \mathfrak{A} \cap S$. Then obviously, $\mathfrak{A}_\lambda \cap S = (\cup_{\alpha<\lambda} \mathfrak{A}_\alpha) \cap S = \cup_{\alpha<\lambda}(\mathfrak{A}_\alpha \cap S) = \mathfrak{A} \cap S$.

(2) Since the set $\{B \mid B \prec_\propto A\}$ is $\propto$-closed, it follows from (1) that

$$\mathfrak{A} \cap \{B \mid B \prec_\propto A\} = lim(\mathcal{N}) \cap \{B \mid B \prec_\propto A\}$$

Also, it holds that $\mathfrak{A} \subseteq lim(\mathcal{N})$, hence

$$\mathfrak{A} \cap \{B \mid B \propto A\} \subseteq lim(\mathcal{N}) \cap \{B \mid B \propto A\}$$

Since $\propto$ is a monotone dependency, it holds that $\mathfrak{A} \vdash_\Delta A$ entails that $lim(\mathcal{N}) \vdash_\Delta A$. □

Proposition 5.5. *Let $\propto$ be a monotone dependency of Δ in $\mathcal{O}$.*

(a) If a natural induction $\mathcal{N} = \langle \mathfrak{A}_\alpha \rangle_{0 \leq \alpha \leq \beta}$ respects $\propto$ then $\mathcal{N}$ is safe.

(b) If for some $i \geq 0$, $\mathfrak{A}_i$ is saturated on $\{B \mid B \propto A\}$ and $A \notin \mathfrak{A}_i$, A is strictly underivable.

(c) If $\mathfrak{A}_i \vdash_\Delta A$ and $\mathfrak{A}$ is saturated on $\{B \mid B \prec_\propto A\}$, then A is safely derivable in $\mathfrak{A}_i$.

Proof. (a) follows from (c). (b) and (c) are straightforward consequences of Proposition 5.4. □

The above proposition is a formalization of properties of informal inductive definitions that, just like the confluence property of inductive constructions, we may easily take for granted but that are indispensable for practical reasoning. Indeed, they offer a way of deciding membership of certain facts in the defined relation while constructing only a fraction of it. To decide whether A belongs to the defined set, we "tweak" a partial induction process towards deriving A or towards saturation on $\{B \mid B \propto A\}$. If A is derived, then it belongs to the defined set. If the natural induction gets saturated on $\{B \mid B \propto A\}$ and has not derived A, A does not belong to the defined set.

Example 5.6. It hold that $\{Q\} \models \neg P \wedge Q$ and $\{Q\} \not\models \neg(\neg P \wedge Q)$. We can prove both using the following very short non-terminal natural induction of $\Delta_\models$:

$$(\{Q\}, \neg P) \rightarrow (\{Q\}, Q) \rightarrow (\{Q\}, \neg P \wedge Q)$$

Indeed, this natural induction respects the induction order $\prec_\models$ of Example 3.25 and hence, it can be extended to a terminal one that converges to the defined set; hence, the

defined set includes $(\{Q\}, \neg P \wedge Q)$. Also, the limit of this short natural induction is saturated on $\{A \mid A \prec_{\models} (\{Q\}, \neg(\neg P \wedge Q))\}$ and $(\{Q\}, \neg(\neg P \wedge Q))$ is not derivable from it. ▲

Theorem 5.7. *Assume $\propto$ monotonically orders Δ in $\mathcal{O}$. Then terminal natural inductions that respect $\propto$ exist and all converge to the same limit. Moreover, the limit is independent of $\propto$. It is the safely defined structure of Δ in $\mathcal{O}$.*

Proof. Existence follows from Proposition 5.1. Terminal natural inductions that respect $\propto$ are safe-terminal natural inductions and all of them converge to the safely defined structure. This structure does not depend on $\propto$. □

We have shown here that our intuition is right: that natural inductions that delay the derivation of defined atoms until it is safe, converge. Moreover, since safe natural inductions do not depend on the induction order, the set defined by a definition over some induction order does not depend on that order.

So far, ordered and iterated definitions were defined as pairs $(\Delta, \propto)$ of rule sets and suitable induction order. The confluence theorem entitles us to drop $\propto$ from the definition.

Definition 5.8. A rule set Δ is an ordered definition in $\mathcal{O}$ if some $\prec$ strictly orders Δ in $\mathcal{O}$. A rule set Δ is an iterated definition in $\mathcal{O}$ if some $\propto$ monotonically orders Δ in $\mathcal{O}$. For any monotone, ordered or iterated definition Δ in $\mathcal{O}$, the structure defined by Δ in $\mathcal{O}$ is the safely defined structure.

6 Other properties of definitions

Safe natural inductions go faster Natural inductions that respect a suitable $\propto$ are safe. The converse does not hold. The following example shows that safe natural inductions do not necessarily respect the induction order, and that they may derive a fact in far fewer steps than an induction that respects the induction order.

Example 6.1. A (two-step) natural induction of the satisfaction definition $\Delta_{\models}$ that does not respect the induction order:

$\rightarrow Sat(\{P\}, P) \rightarrow \{Sat(\{P\}, P \vee \varphi) \mid \varphi$ a formula over the propositional vocabulary $\xi\}$

Indeed, after the first step $Sat(\{P\}, P)$ is derived in structure $\mathfrak{A}_1$. It is easy to see that $\mathfrak{A}_1 \vdash_\Delta Sat(\{P\}, P \vee \varphi)$ for every φ. In fact, since the rule defining $Sat(\{P\}, P \vee \varphi)$ is monotone, the fact is derivable in each superset of $\mathfrak{A}_1$ and hence, it remains derivable in every natural induction from $\mathfrak{A}_1$. Therefore, each domain atom $Sat(\{P\}, P \vee \varphi)$ is safely derivable from $\mathfrak{A}_1$. Of course, the induction does not respect the induction order. ▲

On the level of informal definitions, the safe natural induction in Example 6.1 probably matches how many of us derive the satisfaction of a disjunction $\mathfrak{A} \models \varphi \vee \psi$: if a disjunct φ is derived to be satisfied, we jump to the conclusion that $\varphi \vee \psi$ is satisfied, even if the value of ψ is still unknown. Strictly speaking, here we are violating the

induction order! It is nevertheless safe. This derivation step is a safe one, and any fact derived during a safe natural induction is correct.

A difficulty of computing safe natural inductions is to determine the safety of a derivation. The following proposition gives a simple but useful criterion.

Proposition 6.2. *If for structure $\mathfrak{A}$ and defined domain atom A it holds that $\mathfrak{A}' \vdash_\Delta A$ for every $\mathfrak{A}' \geq_{\mathrm{t}} \mathfrak{A}$, then A is safely derivable from $\mathfrak{A}$.*

Proof. Obvious, since the limit of every natural induction from $\mathfrak{A}$ is a superset of $\mathfrak{A}$. □

The proposition reveals a simple but useful criterion to decide whether a defined atom A is safely derivable at some stage $\mathfrak{A}_i$ of the induction: it suffices that it is deriable by a monotone rule at that stage. This was exploited in Example 6.1. As such the criterion for deciding safety of a derivation at the formal level proven in the proposition explains a common and useful practice with informal definitions.

Natural inductions that respect $\propto$ follow $\propto$ An expected property is that inductions *follow* the induction order, that is, if $B \prec_\propto A$ are both derived by $\mathcal{N}$, then B is derived before A. We already discussed this for ordered inductive definitions. The property holds more generally for iterated inductive definitions.

Proposition 6.3. *Assume that $\propto$ monotonically orders Δ in $\mathcal{O}$. If a natural induction $\mathcal{N}$ respects $\propto$ then $\mathcal{N}$ follows $\propto$.*

Proof. If $A \in \mathfrak{A}_{i+1} \setminus \mathfrak{A}_i$, then $\mathfrak{A}_i$ is saturated in $\{B \mid B \prec_\propto A\}$. Hence, by Proposition 5.4, for every $j \geq i$, $\mathfrak{A}_j|_{\prec_\propto A} = \mathfrak{A}_i|_{\prec_\propto A}$. Hence, if $B \prec_\propto A$ and $B \in \mathfrak{A}_{j+1} \setminus \mathfrak{A}_j$, then it holds that $j < i$. □

Safe natural inductions do not necessarily follow the induction order. E.g., a safe induction that derives first $\{P\} \models P$, then $\{P\} \models P \vee \varphi$ may be extended to derive subformulas of φ and in that case, it does not follow the induction order.

The defined structure is a fixpoint of Δ Another intuitively obvious property of the considered sorts of informal definitions is that the defined set is a fixpoint of its induced operator. Formally, a defined domain atom $P(\bar{a})$ holds if and only if it is derivable by one of the rules. In other words, the defined set is a fixpoint of the induced operator $\Gamma^{\mathcal{O}}_\Delta$ (Definition 3.17).

Proposition 6.4. *Let Δ be a definition by iterated induction over $\propto$ in $\mathcal{O}$. If $\mathfrak{A}$ is the structure defined by Δ in $\mathcal{O}$, then $\mathfrak{A}$ is a fixpoint of $\Gamma^{\mathcal{O}}_\Delta$.*

Proof. Select an arbitrary terminal natural induction $\mathcal{N}$ with limit $\mathfrak{A}$. Then any $A \in \mathfrak{A}$ was safely derived at stage $\|A\|_\mathcal{N}$, and by safety $\mathfrak{A} \vdash_\Delta A$. Vice versa, if $\mathfrak{A} \vdash_\Delta A$ then since $\mathfrak{A}$ is saturated, $A \in \mathfrak{A}$. □

It is well known that the converse property does not hold, i.e., not every fixpoint of the operator is the defined structure of Δ in $\mathcal{O}$. For instance, for any context structure

$\mathcal{O}$, the definition Δ_{TC} of transitive closure has always a fixpoint $\mathfrak{A}$ in which $R^{\mathfrak{A}}$ is the complete binary relation on the domain. Also the iterated inductive definition of satisfaction of multi-agent logic in Example 3.52 has multiple fixpoints.[3]

However, if Δ is an ordered definition by induction over $\prec$, then the fixpoint is unique.

Proposition 6.5. *Let Δ be a definition by ordered induction over the well-founded order $\prec$ in $\mathcal{O}$. $\mathfrak{A}$ is the defined structure by Δ in $\mathcal{O}$ if and only if $\mathfrak{A}$ is a fixpoint of $\Gamma^{\mathcal{O}}_{\Delta}$.*

Proof. If suffices to prove that the operator has only one fixpoint. Assume it has two different fixpoints $\mathfrak{A}, \mathfrak{B}$ and assume that A is a domain atom on which $\mathfrak{A}, \mathfrak{B}$ disagree that is minimal in $\prec$. Such a minimal atom certainly exists since $\prec$ is a strict well-founded order. Then it holds that $\mathfrak{A} \cap \{B \mid B \prec\} = \mathfrak{B} \cap \{B \mid B \prec A\}$. Since $\prec$ is a dependency of Δ, it holds that $\mathfrak{A} \vdash_{\Delta} A$ iff $\mathfrak{B} \vdash_{\Delta} A$. Contraction. □

Proposition 6.6. *Let Δ be a definition by iterated induction over $\propto$ in $\mathcal{O}$. The structure $\mathfrak{A}$ defined by Δ in $\mathcal{O}$ is a minimal fixpoint of $\Gamma^{\mathcal{O}}_{\Delta}$.*

Note that $\Gamma^{\mathcal{O}}_{\Delta}$ may have many minimal fixpoints.

Proof. Assume towards contradiction that the defined structure $\mathfrak{A}$ is not a minimal fixpoint of $\Gamma^{\mathcal{O}}_{\Delta}$ and that $\mathfrak{A}'$ is a strictly lesser one. Consider a natural induction $\mathcal{N}$ that respects $\propto$ and constructs $\mathfrak{A}$. Let i be the minimal ordinal such that some atom $A \in \mathfrak{A} \setminus \mathfrak{A}'$ is derived at $\mathfrak{A}_i$. Since $\mathcal{N}$ respects $\propto$, $\mathfrak{A}_i$ is saturated on the downward closed set $\{B \mid B \prec_{\propto} A\}$. By minimality of i, $\mathfrak{A}_i \subseteq \mathfrak{A}'$ and $\mathfrak{A}_i \cap \{B \mid B \prec_{\propto} A\} = \mathfrak{A} \cap \{B \mid B \prec_{\propto} A\} = \mathfrak{A}' \cap \{B \mid B \prec_{\propto} A\}$. Since $\mathfrak{A}_i \vdash_{\Delta} A$ and $\propto$ is a monotone dependency, it follows that $\mathfrak{A}' \vdash_{\Delta} A$ and hence, that $A \in \Gamma^{\mathcal{O}}_{\Delta}(\mathfrak{A}') = \mathfrak{A}'$. Contradiction. □

Another good question is whether the safely defined set of a formal definition Δ in $\mathcal{O}$ is a fixpoint. It obviously is if the definition is an iterated definition but what if it is not? It will be considered later in this text, when we consider less sensible definitions and definitional paradoxes.

The two experiments We finalize our discussion of the two "experiments" introduced in Section 1.

Example 6.7. In previous examples, we verified that Δ_{TC} faithfully expresses Definition 1.1 and that an induction process of the informal definition in the context of a graph $\mathcal{G}$ corresponds to a natural induction in the corresponding context structure $\mathcal{O}$. It follows that the formally and informally defined sets correspond. ▲

Example 6.8. Likewise, we verified that $\Delta_{\models}$ faithfully expresses Definition 1.2, that a context structure $\mathcal{O}$ corresponds to the pair of sets of structures and formulas for propositional vocabulary ξ, that the induction order $\prec_{\models}$ in context $\mathcal{O}$ corresponds to the subformula order, that natural inductions respecting $\prec_{\models}$ in context $\mathcal{O}$ correspond to induction processes "along" the subformula order. It follows that the formally and informally defined sets correspond. ▲

[3] Fixpoints exist in which all formulas $DC_g\varphi$ are satisfied; the monotone inductive rule maintains this set when applying the operator.

Complexity of computing the safely defined structure in finite context In this section, we investigate the data complexity of deciding a defined domain atom in the safely defined structure. Hence, here we will focus on finite definitions and contexts.

Let Δ be a fixed definition over $\Sigma = pars(\Delta) \cup def(\Delta)$. We assume (without loss of generality due to the finiteness of Δ) that each defined predicate $P \in def(\Delta)$ is defined by exactly one rule

$$\forall \bar{x}(P(\bar{x}) \leftarrow \varphi_P[\bar{x}]).$$

A key computation step in a safe natural induction is the verification that a domain atom $P(\bar{a})$ is safely derivable from $\mathfrak{A}$ in $\mathcal{O}$. We will show that problem is in co-NP.

Theorem 6.9. *For a given finite definition Δ, the problem of deciding whether a defined domain atom $P(\bar{a})$ is safely derivable from a $def(\Delta)$-structure $\mathfrak{A}$ in finite context $\mathcal{O}$ is in co-NP.*

Proof. Algorithm 1 contains a nondeterministic program to decide that $P(\bar{a})$ is *not* safely derivable from $\mathfrak{A}$ in $\mathcal{O}$. It takes as input $\mathfrak{A}, \mathcal{O}, P(\bar{a})$ and Δ. Let D be the domain of $\mathfrak{A}$ and $\mathcal{O}$. This program nondeterministically traverses a natural induction from Δ

Algorithm 1 Nondeterministic algorithm to decide that $P(\bar{a})$ is *not* safely derivable from $\mathfrak{A}$ in $\mathcal{O}$.

while `true` **do**
  $S \leftarrow \left\{ Q(\bar{b}) \in At_D^{def(\Delta)} \mid \mathfrak{A} \models \varphi_Q[\bar{b}] \right\}$
  if $P(\bar{a}) \notin S$ **then**
    return `true`
  else if $S \subseteq \mathfrak{A}$ **then**
    return `false`
  else
    choose a non-empty subset $S' \subseteq S \setminus \mathfrak{A}$
    $\mathfrak{A} \leftarrow \mathfrak{A} \cup S'$
  end if
end while

in $\mathcal{O}$. Every state $\mathfrak{A}'$ that can be reached by a natural induction from $\mathfrak{A}$ can be reached by a run of this program. The algorithm stops with `true` when it reaches a structure in which $P(\bar{a})$ is not derivable; this shows that $P(\bar{a})$ is not safely derivable. It stops with `false` if the reached structure is saturated and still derives $P(\bar{a})$. In this case, the traversed natural induction was a terminal one that was not a witness that $P(\bar{a})$ was not safely derivable.

One run of the algorithm builds a strictly growing sequences of $def(\Delta)$-structures; hence, the number of iterations is bound by the cardinality of the set $At_D^{def(\Delta)}$ of defined domain atoms. This number is polynomial in the size of $\mathcal{O}$. At each step, the main computation is the computation of S, the set of derivable domain atoms from $\mathfrak{A}$. This is a polynomial operation. Thus, this is a nondeterministic polynomial program that has a run that terminates with `true` if and only $P(\bar{a})$ is not safely derivable from $\mathfrak{A}$ in $\mathcal{O}$. It follows that deciding that $P(\bar{a})$ is not safely derivable from $\mathfrak{A}$ in $\mathcal{O}$ is in NP and its dual is in co-NP. □

Theorem 6.10. *For every Δ, the problem of deciding whether a defined domain atom $P(\bar{a})$ holds in the safely defined structure in a finite $pars(\Delta)$-structure $\mathcal{O}$ is in Δ_2^P. For some Δ, the problem is co-NP hard.*

Proof. We first show containment in Δ_2^P. By solving a polynomial number of co-NP problems, we can compute the set of safely derivable atoms in any structure. By doing this a polynomial number of times, we find the safely defined structure and can determine if $P(\bar{a})$ is true in it.

We now show co-NP-hardness. To do this, we encode the co-NP-hard problem of deciding validity of a propositional formula φ in Disjunctive Normal Form (DNF). A formula φ in DNF is a set of disjuncts, each of which is a conjunction of literals. We encode φ over propositional vocabulary ξ as a context structure $\mathcal{O}_\varphi$ providing interpretation for the following parameter predicates:

- domain of $\mathcal{O}_\varphi = \xi \cup \varphi$: it contains all propositional symbols of ξ and all disjuncts of φ.
- $Prop^{\mathcal{O}_\varphi} = \xi$,
- $Dis^{\mathcal{O}_\varphi} = \varphi$, the set of disjuncts of φ;
- $Pos^{\mathcal{O}_\varphi} = \{(d, p) \mid p \text{ is a positive literal in a disjunct } d \in \varphi\}$;
- $Neg^{\mathcal{O}_\varphi} = \{(d, p) \mid \neg p \text{ is a negative literal in a disjunct } d \in \varphi\}$.

The size of $\mathcal{O}_\varphi$ is linear in the size of φ. Furthermore, we introduce the predicate $T/1$ to encode ξ-structures I. In particular, the value of T will be the set of true propositional symbols in I. Consider the following definition:

$$\Delta = \left\{ \begin{array}{l} Val \leftarrow \exists d(Dis(d) \wedge \forall p(Pos(d,p) \Rightarrow T(p)) \wedge \forall p(Neg(d,p) \Rightarrow \neg T(p))) \\ \forall p(T(p) \leftarrow Val \wedge Prop(p)) \end{array} \right\}$$

The definition defines $Val/0$ and $T/1$. It is straightforward to see that the definiens of Val expresses the satisfaction of φ in the ξ-structure encoded by $T/1$.

We prove that Val is true in the safely defined structure in $\mathcal{O}_\varphi$ if and only if the encoded formula φ is valid.

First, if φ is false in $\emptyset$, then φ is not valid. In this case, the first element $\mathfrak{A}_0 = \emptyset$ of every natural induction encodes the empty ξ-structure in T, and we see that $\langle \mathfrak{A}_0 \rangle$ is the uniqe natural induction. Hence, in this case, Val is false in the safely derived structure.

Otherwise φ is true in $\emptyset$. Then Val is derivable from $\mathfrak{A}_0$ but not necessarily safely derivable. In fact, Val is safely derivable exactly if φ is valid. Indeed, the natural inductions of this definition are of the form

$$\rightarrow Val \rightarrow T1 \rightarrow T2 \rightarrow \ldots$$

where each T_i is a set of domain atoms of the form $T(p)$. If φ is valid, then the definiens of Val continues to hold at each stage. On the other hand, if φ is not valid, then there exists a ξ-structure I in which φ is false. Then I is not empty since φ is true in $\emptyset$. Consider the following two step natural induction:

$$\rightarrow Val \rightarrow \{T(p) \mid p \in I\}$$

This is a natural induction obtained by applying the second rule for each $p \in I$ at the second stage. In its limit, the definiens of Val is false and Val is not derivable. Hence, the derivation of Val is not safe.

We conclude that the safely defined structure contains Val if and only if φ is valid. Since deciding validity of a sentence in DNF is co-NP hard, we obtain the desired result. □

The complexity of computing the safely defined structure may be too high for practical computing. Here an important computational use of the induction order surfaces. Given an induction order $\prec$, a defined domain atom A can be decided with a natural induction in at most $\#\{B \mid B \prec A\} + 1$ steps. If rule application is computationally feasible, then this may be an efficient method. Even better, the method may work also in infinite structures. E.g. for Definition 1.2 of the satisfaction relation and its formalization $\Delta_{\models}$, the context structures $\mathcal{O}$ are infinite since they contain infinitely many formulas. Yet, applying a rule is a constant time operation (given the satisfaction of component formulas), and the complexity of deciding the defined domain atom "φ is true in I", is linear in $\#\{\psi \mid \psi \prec_{\models} \varphi\}$, the size of the formula.

Nevertheless, it would be nice to have an inductive construction method that does not require an explicit induction order and that is more efficient (preferably tractable) than computing safe natural inductions. If the polynomial hierarchy does not collapse, it follows from the above theorem that such a method would not always compute the safely defined structure. But perhaps it would be strong enough to compute the defined set of definitions that occur in practice. In an earlier version of this work (Denecker and Vennekens, 2014), we proposed an alternative construction, namely ultimate well-founded induction. However, this suffers from the same complexity problems. We conjecture, and intend to prove in future work, that standard well-founded inductions (Denecker and Vennekens, 2007) provide such a construction: a tractable induction process that is strong enough to compute the defined set of definitions that occur in practice.

Summary: implications for informal definitions This section and the previous one reveal some more properties of the considered types of informal definitions.

First, it confirmed what has been stated at the end of Section 3: that the role of the induction order in an informal definition is to delay rule application till it is safe to do; that safety ensures confluence; that the choice of the induction order does not matter as long as it monotonically orders the definition.

Second, that natural inductions that respect a suitable induction order $\propto$ follow this order: larger atoms in the induction order are derived later.

Third, that in practical dealing with informal definitions, we often do not respect or follow an induction order in the induction process. This is not needed as long as the derivations that we make are safe. In this context, we have seen that a monotone rule can be applied safely at all times.

Fourth, the properties of this type of definitions make it possible to perform many computations in a cheap way, without computing the full induction process. By tweaking the induction process in the right direction, we may efficiently decide membership

of a defined fact. From a pragmatical point of view, this is certainly of crucial importance to reason on informal definitions.

Fifth, the defined set is a fixpoint of the induced operator, and in case of an ordered definition, it is the unique fixpoint. In case of an iterated inductive definition, it is a minimal fixpoint.

As said before, it seems that we often take these properties of informal definitions for granted. The contribution here is that we are able to mathematically prove them in a formal study, often for the first time.

The use of safe natural inductions to construct the defined set has some great potential: they do not require knowledge of an induction order, and they are faster than inductions that respect an induction order. From a knowledge representation perspective, the benefit is that in representing an informal definition, it suffices to represent the rules; there is no need to express the induction order. From a computational perspective, it may be useful that safe inductions are much "faster" and derive defined atoms with far less derivations.

On the other hand, it is clear as well that computing safe inductions could be awfully difficult. Verifying that a fact A is safely derivable in structure $\mathfrak{A}$ seems to require some form of unbounded "lookahead" to verify that it remains derivable in any natural induction from $\mathfrak{A}$. We analyzed the finite case and showed that this check can be co-NP hard. Computing natural inductions respecting a given induction order $\prec$ might be much cheaper, even in infinite structures.

While the induction order does not affect the defined set, it gives insight in the definition, it shows how to set up the induction process towards a query, and provides us with a test to verify the mathematical sensibility of the definition (see the next section). Computationally, it may suggest an efficient method to compute defined facts. Thus indeed, specifying the induction order is useful.

7 Beyond iterated definitions: white, black and different shades of grey

The previous sections present a formal model of monotone, ordered and iterated definitions and define the notion of safely defined structure of a rule set Δ in context $\mathcal{O}$. Most formal science theories are approximations of the studied reality and our theory is no exception. There is a core area of informal definitions that the theory captures perfectly, but there are border cases as well. It is important to develop an understanding of these border cases and how the theory behaves on them.

Sensible definitions beyond iterated definitions The informal semantics of rule sets as definitions does not abruptly break beyond iterated definitions, as shown by the following example.

Example 7.1. Let $\mathcal{O}$ be the natural number context structure of Example 3.11 and Δ_{ev1} the following variant definition of Δ_{ev} that defines both $Even$ and an auxiliary

predicate $Next$:

$$\left\{ \begin{array}{l} \forall x \forall y (Next(x,y) \leftarrow x = y+1) \\ \forall x (Even(x) \leftarrow x = 0 \vee \exists y (Next(x,y) \wedge \neg Even(y))) \end{array} \right\}$$

This rule set faithfully expresses what to us seems an acceptable mathematical definition of the set of even numbers in the context of the natural numbers. The safely defined model $\mathfrak{A}$ of this definition in $\mathcal{O}$ is the intended one. Indeed, the following natural induction is safe.

$$\rightarrow \{Next(n, n+1) \mid n \in \mathbb{N}\} \rightarrow Ev(0) \rightarrow Ev(2) \rightarrow \ldots$$

Nevertheless, this rule set is not an ordered or iterated inductive definition in $\mathcal{O}$ according to Definition 5.8. Indeed, in any dependency relation of Δ_{ev1}, it holds that $Ev(n) \propto Ev(m)$ for all $n, m \in \mathbb{N}$. This follows from the fact that $Ev(m)$ is derivable in the structure $\{Next(m,n)\}$ but not in $\{Next(m,n), Ev(n)\}$. However, as the same two structures show, no such $\propto$ monotonically orders Δ_{ev1} in $\mathcal{O}$. ▲

By all means, the definition Δ_{ev1} is an innocent syntactic variation of Δ_{ev}. It was obtained by applying a general, useful technique: expliciting the definition of an intermediate concept in a definition. The example shows that this operation may easily break an iterated inductive definition. As such, the above example shows a disturbing brittleness of the concept of a definition by iterated induction as defined in Definition 3.46 that fortunately is not shared by the rule formalism under the semantics of safely defined structures.

(Partially) paradoxical definitions Not every (informal or formal) rule set is a sound mathematical definition. E.g., "we define a natural number to be Foo if it is not Foo". It is faithfully expressed as:

$$\Delta_{Foo} = \left\{ \ \forall x (Foo(x) \leftarrow \neg Foo(x)) \ \right\}$$

Intuitively, this is a paradoxical definition: if some number is not Foo, then it is per definition Foo, but if it is Foo, then it is per definition not Foo. On the other hand, the safely defined set is well-defined in every context $\mathcal{O}$: it is the empty set (the natural induction $\langle \mathfrak{A}_0 \rangle$ is safe-terminal). Every domain atom is derivable in the safely defined set but none is safely derivable. Hence, the safely defined set is not saturated, it is not a fixpoint nor does it satisfy the implications of the definition.

Other rule sets do not have that paradoxical flavour of the Foo definition but nevertheless fail to define a set. E.g.,

$$\left\{ \begin{array}{l} \forall x (P(x) \leftarrow \neg Q(x)) \\ \forall x (Q(x) \leftarrow \neg P(x)) \end{array} \right\}$$

Also for this definition, the one step natural induction $\langle \mathfrak{A}_0 \rangle$ is safe-terminal and the safely defined structure in every context is the empty structure. Every atom of P and Q is derivable but none is safely derivable.

Some definitions are mathematically sound definitions in some contexts and not in others. But in these other contexts, they still *partially* define a set.

Example 7.2. In the context of the integer numbers $\mathbb{Z}$, the evenness definition Δ_{ev} is not an ordered definition. It has a unique natural induction that is safe-terminal but not terminal.

$$\rightarrow Ev(0) \rightarrow Ev(2) \rightarrow \ldots$$

The safely defined set is the set of positive even numbers. Each atom $Even(n)$ for negative number n is derivable but none is safely derivable. The safely defined structure is not saturated, it is not a fixpoint and it does not satisfy the implications of the definition. ▲

Some informal rule sets (in some context $\mathcal{O}$) are not sound definitions according to mathematical standards. It seems to us that such definitions would be considered as mathematical errors and they should not be allowed to appear in reviewed mathematical text. They may still *partially* define a set: some objects are soundly derived to be in the set, others to be out the set, and some are undecided.

On the other hand, the safely defined structure is a mathematically well-defined structure for every definition Δ in every context $\mathcal{O}$. One clear indication of an error is when the safely defined structure is not saturated. The derivable but not safely derivable domain atoms are undecided elements of the defined set. Also atoms that are not derivable but not strictly underivable are undecided elements of the defined set.

While such (partially) paradoxical definitions seem unacceptable in standard mathematical practice, certain *definitional paradoxes* and the partial sets they define have attracted considerable attention in the philosophical logic community.

Theory of truth A longstanding problem in philosophical logic is the definition of a truth predicate (Tarski, 1944; Kripke, 1975). The exposition below is based on (Fitting, 1997). Let Σ be the vocabulary of Peano arithmetic (possibly augmented with additional symbols), and $\mathcal{O}$ its standard interpretation (extended for the additional symbols). Let $T/1$ be a new unary predicate and $\Sigma_T = \Sigma \cup \{T/1\}$. Assume there is a Gödel numbering $\lceil\cdot\rceil$ of formulas over Σ_T which allows for paradoxes and other self-referential statements such as *liars* "this sentence is false" and *truth tellers* "this sentence is true". The challenge is to define T as the set of all Gödel numbers of true propositions. Formally, its (infinite) definition Δ_{Truth} consists of, for each sentence φ over Σ_T, the rule

$$T(\lceil\varphi\rceil) \leftarrow \varphi$$

This is a recursive definition since sentences φ are allowed to include the truth predicate. A liar sentence has the form $\neg T(n_l)$ where $n_l = \lceil\neg T(n_l)\rceil$. For such a liar sentence, the definition contains the rule:

$$T(n_l) \leftarrow \neg T(n_l)$$

Thus, $T(n_l)$ is undefined in the defined structure of Δ_{Truth}. A truth teller has the form $T(n_t)$ where $n_t = \lceil T(n_t)\rceil$. For them, Δ_{Truth} contains rules of the form:

$$T(n_t) \leftarrow T(n_t)$$

Kripke proposed a three-valued construction that produces a partial set for T in which liar and truth sayer sentences are left undefined, as well as many other self-referential sentences.

We can see that in the safely defined structure of this definition, liars and truth sayers are false. Liars are not safely derivable, nor strictly underivable, while truth sayers are strictly underivable.

Not so sensible monotone, ordered, iterated definitions A key intuition that has guided the research here is the idea that a rule can be applied only if it continues to apply in every natural induction. The concept of safe natural induction formalizes this by only deriving atoms A in structure $\mathfrak{A}$ that remain derivable in every natural induction from $\mathfrak{A}$. However, this is not *exactly* the same. In the remainder of this section, we will see examples of monotone, ordered, iterated definitions and safe natural inductions where a rule is "safely applied", and yet, the rule condition becomes violated at a later stage.

Example 7.3. The rule set $\left\{\ P \leftarrow \mathbf{t}\ \right\}$ is a monotone, an ordered and an iterated definition. The induction order is the empty order $\prec_\emptyset$. If we replace its rule body by a tautology, these properties and the defined structure are preserved. (see Proposition 7.4 below). Hence, also

$$\left\{\ P \leftarrow \neg P \vee P\ \right\}$$

and, after splitting this rule:

$$\left\{\begin{array}{l} P \leftarrow \neg P \\ P \leftarrow P \end{array}\right\}$$

are monotone, ordered and iterated definitions that define $\{P\}$. Notice that in the latter rule, P depends on itself in each rule separately, yet globally it does not depend on itself. The natural induction

$$\rightarrow P$$

derives P using the first rule. The condition of this rule is violated after application of the rule. However, by then the second rule applies. Hence, this is a safe-terminal natural induction. ▲

The correctness of the transformation in the above example follows from the proposition below.

Proposition 7.4. *Let φ, φ' be logically equivalent. Substituting rule $\forall\bar{x}(P(\bar{t}) \leftarrow \varphi')$ for a rule $\forall\bar{x}(P(\bar{t}) \leftarrow \varphi)$ in Δ in some context $\mathcal{O}$ preserves (monotone) dependencies, the property of being a monotone/ordered/iterated definition, (safe) natural inductions and the defined set. The same holds for splitting a rule $\forall\bar{x}(P(\bar{t}) \leftarrow \varphi \vee \psi)$ in a pair $\forall\bar{x}(P(\bar{t}) \leftarrow \varphi), \forall\bar{x}(P(\bar{t}) \leftarrow \psi)$.*

Proof. The concepts of dependency, monotone dependency, and (safe) natural induction are defined semantically and hence, they are preserved under equivalence preserving transformations to rule bodies. □

It is questionable whether the behaviour displayed in the previous example is found in informal definitions in mathematical text. At least, we have never seen this. It would occur in the variant informal definition in the following example.

Example 7.5. Consider the variant of the informal definition of satisfaction (Definition 1.2) obtained in a similar way, by replacing its first rule by the following ones:

- $I \models P$ if $I \not\models P$ and $P \in I$;
- $I \models P$ if $I \models P$ and $P \in I$;

One could argue that these new rules "obviously" are equivalent to the original one by appealing to the fact that "$I \models \neg P$ or $I \not\models \neg P$" is tautologically true. But the modified rules intuitively mismatch the induction order and we doubt whether such a definition would be accepted in mathematical text (we would not accept it, at least).

On the formal level, $\Delta_{\models}$ is easily modified to express the above informal definition. It is an ordered definition in every suitable context $\mathcal{O}$ induced by a propositional vocabulary ξ. For $P \in I$, any safe natural induction derives $I \models P$ using the first rule, after which its condition is violated but the second rule starts to apply. ▲

The same behaviour as in the previous case is found also in rule sets that are not iterated definitions.

Example 7.6. The rule set below has the property that its safely defined structure is a non-minimal fixpoint.

$$\left\{ \begin{array}{l} Q \leftarrow \neg P \\ Q \leftarrow P \wedge Q \\ P \leftarrow Q \\ P \leftarrow P \end{array} \right\}$$

It is not an iterated definition : the unique dependency is the total relation on defined atoms, and this does not monotonically order the definition. The unique terminal natural induction of this definition is:

$$\rightarrow Q \rightarrow P$$

Formally, it is safe. Indeed, initially Q is derived from its first rule. While its condition $\neg P$ is later canceled, it is canceled only after the derivation of P, at which time Q is derivable from its second rule. The defined set is a fixpoint. However, it is not a minimal fixpoint. The least fixpoint is $\{P\}$. ▲

In all above cases, we found behavior that we probably never encounter in mathematics. This suggests that an implicit convention of informal definitions is not yet explicit in our theory: our definitions of monotone, ordered and iterated definitions and the concept of safe natural induction might be too liberal. Some convention regarding informal definitions remains to be discovered.

8 On the role of the induction order

The role of the induction order $\prec$ in an informal definition is to delay the application of a (nonmonotone) rule until it is *safe* to do so, as explained above. This ensures that all natural inductions that respect $\prec$ are confluent (if at least the induction order is a dependency of the informal definition). But contrary to what might be expected, the induction order has no semantical role: the set defined by a definition over an induction order can be computed without knowing this order and hence, it does not depend on the induction order.

Nevertheless, there is value in specifying an induction order for (informal) definitions.

- An induction order gives the reader insight in the structure of the definition, and the structure of its (safe or natural) induction processes. It helps the reader to understand the definition.
- Given that it is not difficult to write senseless definitions, specifying an induction order yields a "parity check" for verifying the mathematical sensibility of the definition. Once the reader has verified that the induction order is a well-founded order and dependency of the definition, she can be certain that the definiendum is well-defined. One will "feel" this parity check in operation when reading the following toy definition:

  Definition 8.1. We define the set of grue natural numbers by induction on the standard order:

  - n is grue if $n + 1$ is grue.

  This is a monotone definition defining the set of grue numbers to be empty, but the induction order does not match the inductive rule, and such a definition is mathematically unacceptable.
- From a reasoning and computational point of view, a given induction order is a tremendous help. It helps to build safe natural inductions without the expensive "safety check".
- It helps to tune an induction process to compute truth or falsity of a defined fact.
- The induction order gives us a criterion to stop a partial induction process and still be certain that a queried fact cannot be derived anymore.
- The induction order is underlying all top-down computation procedures to compute the truth value of a defined fact A. Such computations can be understood as intertwined computations of the set $\{B \mid B \prec A\}$ and the computation of a natural induction on this set towards A.
- In infinite spaces, terminal natural inductions are infinite objects, but if $\{B \mid B \prec A\}$ is finite, the value of a defined fact A can be computed, and the size of $\{B \mid B \prec A\}$ is an indicator of the computational complexity of this.

Thus, although the induction order $\prec$ (and by extension, an induction relation $\propto$ for iterated definitions) is semantically not useful and it may put a burden on the knowledge representation, it might still be valuable to extend a definition logic to allow to express the induction order. Such an explicit induction order could then be used by the system for the various tasks that we discussed above.

9 Conclusion

This paper could be viewed as an empirical, formal, exact scientific study of certain classes of informal definitions, in the following sense of these words: empirical: definitions exist, can be written, read, interpreted, reasoned upon; formal: a mathematical model is built for them; exact: the formal model characterises the informal definitions in detail.

For lecturers and authors of text books in mathematics and logic, the results of this paper may in the long run prove useful to explain students the meaning of various types of definitions that appear in their courses or text books.

To metamathematics, the paper makes a contribution of formalizing the induction process in common forms of definitions. It led us to explore several fundamental aspects that were not studied before. To recall the most important ones: the non-determinism of the induction process, the role of the induction order and its link with the rules, the confluence of induction processes, the independency of the defined set of the induction order. We pointed to the pragmatical importance of these properties, e.g., the possibility of directing the induction process towards answering a specific query.

The independency of the induction order suggests to build a general rule-based definition logic without induction order, using the semantics of safely defined structures. Our exploration of this idea had a mixed outcome with on the positive side sensible definitions beyond iterated definitions and the link with definitional paradoxes but on the negative side the high complexity of the semantics in finite structures, a lack of distinction between atoms that are defined as false and that are undefined, and also some counterintuitive examples with, e.g., non-minimal defined fixpoints. The paper is open-ended on this level. We conjecture that methods of three-valued logic inspired by the semantics of logic programming are useful to tackle at least some of these problems.

To knowledge representation, the paper contributes a study of an important form of human knowledge available in virtually all KR applications, and supported in many expressive declarative systems.

From the perspective of KR, an uncommon aspect of our study is its exact, formal approach to the studied forms of knowledge. KR research aims to develop formal languages and methods to express knowledge, but it rarely pretends to be the scientific study of knowledge. Knowledge [4] is a cognitive reality that many consider to be out

[4]The word "knowledge" is not he appropriate term here, but the right word does not seem to exist yet, or at least, we do not know it. In philosophical logic, knowledge is often defined as "true justified belief", and this is certainly not what we have in mind. What we have in mind with is the "thing" that could be true or false, believed or not believed, and if it is believed, could be believed in a justified way or in an unjustified way. It could perhaps better be called "a piece of information", or a "quantum of information", a word used in (Devlin, 1991) in a sense that is certainly close that what we had in mind.

of reach of the methods of formal empirical scientific investigation. Moreover, human knowledge is communicated in informal natural language which for many is inevitably vague and ambiguous. It is true that natural language is sometimes ambiguous, but there are other contexts where it can be extremely precise. In particular, mathematicians and formal scientist are trained in precise language and they use informal natural language to build their disciplines with mathematical precision. The definition is a part of the informal language of formal science and mathematics. Its precision makes it an ideal target for a formal 'empirical' study. In this, we have stressed the mathematical nature of the "experiments". E.g., the question of whether an informal definition such as Definition 1.2 is faithfully expressed by a formal definition such as $\Delta_{\models}$, whether its subformula order is formalized by $\prec_{\models}$, or whether the formally and informally defined sets coincide: these facts are not matter of contention; they are mathematical facts.

A contribution that is not yet elaborated here is to logic programming. There is an obvious syntactical link between definitions in this papers and logic programs. Ever since the early days of logic programming when the negation as failure problem arose, some have explained the declarative meaning of logic programs in terms of definitions. This holds especially for those that adopted the well-founded semantics (Van Gelder et al., 1991). The link was made explicit in several publications (Schlipf, 1995a; Denecker, 1998; Denecker et al., 2001; Denecker and Ternovska, 2008), where various arguments can be found that, under the well-founded semantics, each rule in a set of rules can be seen as an (inductive or base) case of a (possibly inductive) definition. A problem with all these attempts has been the absence of a natural formal model of the relevant sorts of informal definitions. Therefore, a weakness in all these studies is the absence of an obvious connection between how we understand various types of informal inductive definitions in mathematical text and the complex mathematics of the well-founded semantics. With the present paper, we believe to have closed this gap. Now, we are on mathematical ground. We conjecture that for all mathematical definitions that occur in practice, the set defined by a formal definition is the well-founded model. It is part of our future research agenda to *prove* that this indeed holds.

References

Serge Abiteboul, Richard Hull, and Victor Vianu. 1995. *Foundations of Databases*. Addison-Wesley.

Peter Aczel. 1977. An Introduction to Inductive Definitions. In *Handbook of Mathematical Logic*, J. Barwise (Ed.). North-Holland Publishing Company, 739–782.

Krzysztof R. Apt, Howard A. Blair, and Adrian Walker. 1988. Towards a Theory of Declarative Knowledge. In *Foundations of Deductive Databases and Logic Programming*, Jack Minker (Ed.). Morgan Kaufmann, 89–148.

Ronald J. Brachman, Richard E. Fikes, and Hector J. Levesque. 1983. Krypton: Integrating Terminology and Assertion. In *National Conference on Artificial Intelligence(AAAI'83)*. William Kaufman, 31–35.

Julian C. Bradfield. 1996. The Modal μ-calculus Alternation Hierarchy is Strict. In *CONCUR (LNCS)*, Ugo Montanari and Vladimiro Sassone (Eds.), Vol. 1119. Springer, 233–246.

Wilfried Buchholz, Solomon Feferman, Wolfram Pohlers, and Wilfried Sieg. 1981. *Iterated Inductive Definitions and Subsystems of Analysis: Recent Proof-Theoretical Studies*. Lecture Notes in Mathematics, Vol. 897. Springer.

Keith L. Clark. 1978. Negation as Failure. In *Logic and Data Bases*. Plenum Press, 293–322.

Broes De Cat, Bart Bogaerts, Maurice Bruynooghe, Gerda Janssens, and Marc Denecker. 2016. Predicate Logic as a Modelling Language: The IDP System. *CoRR* abs/1401.6312v2 (2016). `http://arxiv.org/abs/1401.6312v2`

Marc Denecker. 1998. The Well-Founded Semantics Is the Principle of Inductive Definition. In *JELIA (LNCS)*, Jürgen Dix, Luis Fariñas del Cerro, and Ulrich Furbach (Eds.), Vol. 1489. Springer, 1–16.

Marc Denecker, Maurice Bruynooghe, and Victor Marek. 2001. Logic programming revisited: Logic programs as inductive definitions. *ACM Trans. Comput. Log.* 2, 4 (2001), 623–654.

Marc Denecker and Eugenia Ternovska. 2007. Inductive Situation Calculus. *Artif. Intell.* 171, 5-6 (2007), 332–360.

Marc Denecker and Eugenia Ternovska. 2008. A Logic of Nonmonotone Inductive Definitions. *ACM Trans. Comput. Log.* 9, 2, Article 14 (April 2008), 52 pages. `http://dx.doi.org/10.1145/1342991.1342998`

Marc Denecker and Joost Vennekens. 2007. Well-Founded Semantics and the Algebraic Theory of Non-monotone Inductive Definitions. In *LPNMR (Lecture Notes in Computer Science)*, Chitta Baral, Gerhard Brewka, and John S. Schlipf (Eds.), Vol. 4483. Springer, 84–96. `DOI: http://dx.doi.org/10.1007/978-3-540-72200-7_9`

Marc Denecker and Joost Vennekens. 2014. The Well-Founded Semantics Is the Principle of Inductive Definition, Revisited. In *KR*, Chitta Baral, Giuseppe De Giacomo, and Thomas Eiter (Eds.). AAAI Press, 1–10. http://www.aaai.org/ocs/index.php/KR/KR14/paper/view/7957

Keith Devlin. 1991. *Logic and information*. Cambridge University Press.

Solomon Feferman. 1970. Formal theories for transfinite iterations of generalised inductive definitions and some subsystems of analysis. In *Intuitionism and Proof theory*, A. Kino, J. Myhill, and R.E. Vesley (Eds.). North Holland, 303–326.

Melvin Fitting. 1997. A theory of truth that prefers falsehood. *Journal of Philosophical Logic* 26 (1997), 477–500.

Yuri Gurevich and Saharon Shelah. 1985. Fixed-Point Extensions of First-Order Logic. In *26th Annual Symposium on Foundations of Computer Science, Portland, Oregon, USA, 21-23 October 1985*. IEEE Computer Society, 346–353. DOI: http://dx.doi.org/10.1109/SFCS.1985.27

Lars Hallnäs. 1991. Partial Inductive Definitions. *Theor. Comput. Sci.* 87, 1 (1991), 115–142. DOI: http://dx.doi.org/10.1016/S0304-3975(06)80007-1

Ping Hou. 2010. *Techniques for Reasoning in FO(ID) and Fixpoint Logic*. Ph.D. Dissertation. Department of Computer Science, K.U.Leuven, Leuven, Belgium.

Dexter Kozen. 1983. Results on the Propositional μ-Calculus. *Theoretical Computer Science* 27 (1983), 333–354.

Georg Kreisel. 1963. *Generalized inductive definitions*. Technical Report. Section III in the Stanford University report on the Foundations of Analysis.

Saul Kripke. 1975. Outline of a theory of truth. *Journal of Philosophy* 72 (1975), 690–712. DOI:http://dx.doi.org/10.2307/2024634

Michael Leuschel and Michael J. Butler. 2008. ProB: An automated analysis toolset for the B method. *STTT* 10, 2 (2008), 185–203.

W. Marek. 1989. Stable Theories in Autoepistemic Logic. *Fundamenta Informaticae* 12, 2 (1989), 243–254.

Per Martin-Löf. 1971. Hauptsatz for the intuitionistic theory of iterated inductive definitions. In *Second Scandinavian Logic Symposium*, J.e. Fenstad (Ed.). 179–216.

Yiannis N. Moschovakis. 1974a. *Elementary Induction on Abstract Structures*. North-Holland Publishing Company, Amsterdam- New York.

Yiannis N. Moschovakis. 1974b. On non-monotone inductive definability. *Fundamenta Mathematica* 82 (1974), 39–83.

Nicholas Nethercote, Peter J. Stuckey, Ralph Becket, Sebastian Brand, Gregory J. Duck, and Guido Tack. 2007. MiniZinc: Towards a Standard CP Modelling Language. In *CP'07 (LNCS)*, C. Bessiere (Ed.), Vol. 4741. Springer, 529–543.

Karl Raimund Popper. 1959.

Emil L. Post. 1943. Formal Reductions of the General Combinatorial Decision Problem. *American Journal of Mathematics* 65, 2 (1943), 197–215. dx.doi.org/10.2307/2371809

Davide Sangiorgi. 2009. On the origins of bisimulation and coinduction. *ACM Trans. Program. Lang. Syst.* 31, 4 (2009). DOI: http://dx.doi.org/10.1145/1516507.1516510

John S. Schlipf. 1995a. Complexity and Undecidability Results for Logic Programming. *Annals of Mathematics and Artificial Intelligence* 15, 3-4 (1995), 257–288. http://dx.doi.org/10.1007/BF01536398

John S. Schlipf. 1995b. The Expressive Powers of the Logic Programming Semantics. *J. Comput. Syst. Sci.* 51, 1 (1995), 64–86.

Clifford Spector. 1961. Inductively defined sets of natural numbers. In *Infinitistic Methods (Proc. 1959 Symposium on Foundation of Mathematis in Warsaw)*. Pergamon Press, Oxford, 97–102.

Alfred Tarski. 1944. The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research* 4 (1944).

Allen Van Gelder. 1993. The Alternating Fixpoint of Logic Programs with Negation. *J. Comput. Syst. Sci.* 47, 1 (1993), 185–221.

Allen Van Gelder, Kenneth A. Ross, and John S. Schlipf. 1991. The Well-Founded Semantics for General Logic Programs. *J. ACM* 38, 3 (1991), 620–650. DOI: http://dx.doi.org/10.1145/116825.116838